


McGill Undergraduate Mathematics Journal



the δ elta ϵ psilon

Spring 2026

Ninth Issue



Copyright © *The Delta Epsilon*, 2026

All rights reserved. The text of this publication may not be reproduced in any form or by any means without prior written consent of the authors and of *The Delta Epsilon*.

Issue 9, published April 10th, 2026

ISSN: 1911-8996

LETTER FROM THE EDITORS

You might have thought the wait for the next issue of *The Delta Epsilon* would go on forever, but like a perfect square hiding in plain sight, 9 was always going to be just around the corner!

We are beyond thrilled to present Issue 9, which is a continuation of something that once felt like a distant limit but has now firmly converged. Much like a point-cloud accumulating mass before a downpour of proofs, what started as a small idea quickly grew into a robust publication shaped by dedication, collaboration, and few too many \LaTeX battles. Along the way, we learned to *lizard* out the details, *frog-et* our doubts and leap into new challengers, and to stay grounded even when our thoughts drifted among the *point-clouds* of abstraction. At times, the process felt as stable as *German bonds*, reliable and structured, though with admittedly more compilation errors.

This issue is packed with insightful articles, clever puzzles, and reflections from professors and students alike. Every page is a testament to the time and effort poured into making this something truly special for the McGill mathematical community. Please visit our website to access the full version, including puzzle solutions and complete interviews! If you'd like to be featured in the next issue, we encourage you to submit articles next year or apply to be on our editorial board.

We've worked incredibly hard to bring this to life, and we hope you enjoy reading it as much as we enjoyed creating it.

Sincerely,
Aahaan and Hy, Editors-in-Chief
(On behalf of the editorial team)



LETTER FROM SUMS

Last year's revival of our undergraduate mathematics journal was an extraordinary milestone after a decade-long pause. Since then, I have watched what began as the concerted efforts of a determined few, revitalize our corner of the McGill community. I now overhear students talking about *The Delta Epsilon* beyond the confines of our beloved Burnside. I have seen how this journal encourages the pursuit of mathematical curiosity and fosters partnerships between classmates and professors alike.

This ninth edition highlights the power and importance of community in the pursuit of mathematics. Great mathematicians arise from collaboration, discussion, and enormous amounts of curiosity. *The Delta Epsilon* is proof of that. Each submission is shaped through collaboration and represents a wealth of knowledge being passed forward. Thus, the pages you are about to read represent a celebration of the joy of mathematics. Each article perfectly captures the curiosity that drew us into this field and reflects the spirit of determination and pride that define McGill's undergraduate mathematics community.

On behalf of the Society of Undergraduate Mathematics Students, I would like to extend my deepest gratitude to the entire team behind this issue. I am so glad the launch of your newest issue gives the McGill math community a chance to celebrate the efforts you have sustained all year. Your work helps highlight the efforts of every student who has contributed to these pages and will impact many more who will read them. Congratulations to all those published here and thank you to every undergraduate and graduate student who has helped make it happen. You have built something that I have no doubt will outlast us all.

Sincerely,
Madison Cooper, SUMS President
(On behalf of the SUMS council)



THE δ ELTA ϵ PSILON

MCGILL UNDERGRADUATE MATHEMATICS JOURNAL

CONTENTS

The Average Topology of Point Clouds	6
<i>Axel Liu-Bai</i>	
Statistical Inference for Survival Data	11
<i>Jiajun Zhang</i>	
Recurrence on the Frog Model	19
<i>Angela Ji and Yanshuo Liu</i>	
Interview with Anush Tserunyan	24
<i>Louis Meunier and Nitya Khirwar</i>	
Optimal Scoring Strategies for the Lizard Cult	26
<i>Cedric Phillips and Adelaide Row</i>	
A Special Case of Telescopic Series	33
<i>Sebastien Renard</i>	
Oliver Heaviside: A Biography and Contributions	35
<i>Hy Vu</i>	
Fermat's Last Crossword	39
<i>Max Gross</i>	
Cheeger's Inequality	41
<i>Olivia Choi</i>	
Multivariate Composite Copulas: Theory and Application to German Government Bond Yields	46
<i>Gabriella Chen</i>	
Interview with Robert Robere	59
<i>Aahaan Rawal and Aditya Sharma</i>	
Proof of Riemann Mapping Theorem	61
<i>Nitya Khirwar</i>	
The Minimum-Norm OLS Estimator: Unbiasedness and Optimality in the Overparameterized Regime	68
<i>Margaux Butterfield, Perry Santry, Noé Vartanian</i>	
Puzzles	76
Acknowledgements	78

THE AVERAGE TOPOLOGY OF POINT CLOUDS

Axel Liu-Bai

Given a point cloud, we aim to extract its significant topological features, such as connected components and loops, using topological data analysis (TDA). Persistent homology is a tool in TDA that tracks how those features appear and disappear as we vary the maximal distance between two points to connect them. We will present the main ideas behind persistent homology to give an idea of the computations happening when using the functions in Python, and the different representations of the results, such as barcodes, persistence diagrams, and persistence landscapes. When dealing with multiple samples obtained from the same population, it is not entirely clear how we can extract the “average topology” of the underlying point cloud. With this goal in mind, we introduce the mean landscape and discuss methods to estimate it from our data set. To illustrate the concepts, we will apply these tools to a data set of earthquake epicentres.

1 INTRODUCTION

A **point cloud** is a finite subset of R^k denoted by X_n , where n is the number of elements in the subset. For example, the location of earthquake epicentres can be represented as a point cloud in 3 dimensions (latitude, longitude, and depth) or 2 dimensions (latitude and longitude).

Topology is a branch of mathematics that studies properties of geometric shapes that remain invariant under continuous deformations, and we use it to study the shape of point clouds to differentiate one from another.

For example, Table 1 shows point clouds with clear distinguishing topological features. A connected component is defined as a cluster of points, and a loop is a path with the same initial and end point. We can extract one connected component and one loop in a), one connected component and no loop in b), two connected components and two loops in c), and one connected component and two loops in d).

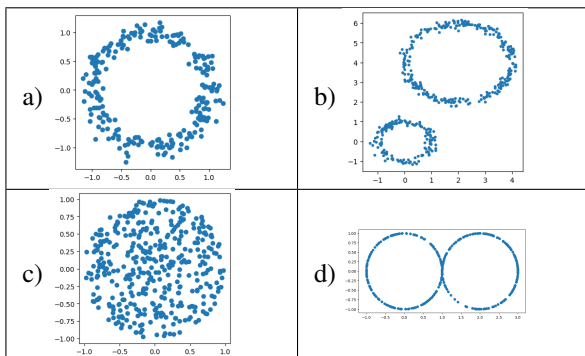


Table 1: Point clouds with significant features

Topological Data Analysis (TDA) aims to detect and quantify the features in point clouds of various complexity, such as multi-dimensional data or samples of larger size. With large sample sizes, TDA computations can take a significant amount of time to run. A

solution is to run the program on smaller samples and then take the average of the detected features. Finding an average of these features is not straightforward. For example, taking the average number of connected components across samples could result in a rational number. Thus, we will present the concept of mean landscape to find the average topology of point clouds.

The example we will be working with to illustrate the concepts of TDA is the data from earthquakes of magnitude 5.5 and above in the region of South America (latitude/longitude rectangle: $[-60, 20] \times [-120, -35]$) from 2000 to 2024 (Figure 1). This data is obtained from <https://earthquake.usgs.gov/earthquakes/search/>. We expect the TDA methods, implemented using functions from the GUDHI Python library [6], to output three loops when applied to this example.




Figure 1: Earthquakes of magnitude 5.5+

2 A PRIMER ON TOPOLOGICAL FEATURES: BETTI NUMBERS

We now formalize this idea of counting connected components and loops. The topological features of a point cloud are captured using an algebraic tool called

homology theory. While a full introduction to homology theory is beyond the scope of this paper, interested readers may refer to *Algebraic Topology* by Hatcher [7].

In essence, homology takes a subset X of \mathbb{R}^k as input and outputs a sequence b_0, b_1, b_2, \dots , of natural numbers called the **Betti numbers of X** . These numbers describe the topological structure of X : b_0 counts the number of pieces making up X (connected components), b_1 the number of loops, b_2 the number of cavities (3-dimensional holes), and so on. In general, b_n represents the number of essential $n + 1$ -dimensional holes. **Essential** means it is not a boundary and it cannot be constructed from other holes. We will illustrate these two conditions in the example below ¹. The intuitive motivation for using Betti numbers is that we can distinguish objects by counting their holes, as illustrated in the introduction.

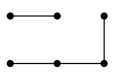
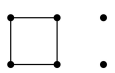
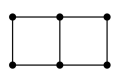

X				
b_0	2	3	1	1
b_1	0	1	2	0

Table 2: Examples of topological spaces and their Betti numbers

We show examples of graphs and their corresponding Betti numbers in Table 2. In the third example, we can count two small loops L_1, L_2 around the squares and a larger loop L around the rectangle as shown in Figure 2. Its $b_1 = 2$ and not 3, since we can construct L from L_1 and L_2 . Intuitively, we can start at the vertex v and go clockwise around L_1 , then clockwise around L_2 . The middle edge disappears as it is crossed twice, but in opposite directions, and the total traversed path is L .




Figure 2: Example 3 of Table 2: Essential and non-essential loops

In the last example of Table 2, the loop around the triangle is not essential since it is now a boundary.

intuitive idea is that since the triangle is included, we can shrink the loop through it. This makes the loop collapse into a point, so we consider the loop non-essential.

The number of loops, represented by b_1 , will be the focus of the mean landscape presented later in this article. It is a distinguishing feature of point clouds and it is less expensive to compute than other Betti numbers since it counts holes of lower dimension ($n + 1 = 2$). Loops were also used to differentiate two different conformations of a protein in the bacteria *E. coli* [8].

3 BASICS OF TOPOLOGICAL DATA ANALYSIS

We cannot directly compute the Betti numbers of our given point cloud of size n because we would get n connected components and no additional information. Instead, we add edges between points that are close together. This is done formally using the **Vietoris-Rips complex** of the point cloud [3]. Suppose we pick a distance α . We build the complex as follows. The vertices of the Vietoris-Rips complex are all the points in the point cloud. Then, we add an edge between vertices x_i, x_j , denoted $[x_i, x_j]$, if $d(x_i, x_j) \leq \alpha$. If we have three vertices, x_i, x_j, x_k , such that all pairwise distances are less than α , *i.e.* they are pairwise connected by edges, then we add a triangle. The boundary of the triangle is precisely those three edges $\{[x_i, x_j], [x_i, x_k], [x_j, x_k]\}$. To generalize this process to higher dimensions, we use simplices. A k -**simplex** $\Delta^k = [x_0, \dots, x_k]$ is in the complex if $d(x_i, x_j) \leq \alpha, \forall i, j \in \{0, \dots, k\}$. In summary, the Vietoris-Rips complex is the collection of all the simplices.

To illustrate how to build the complexes, consider the points $\{(0, 0), (0, 5), (12, 0), (12, 7)\}$, which are the vertices. With $\alpha = 6$ (Figure 3a), we also get a 1-simplex: the edge $[(0, 0), (0, 5)]$. With $\alpha = 13$ (Figure 3b), we get a complex with five 1-simplices, which are the edges, and two 2-simplices, which are the two shaded triangles $\{(0, 0), (0, 5), (12, 0)\}$ and $\{(0, 5), (12, 0), (12, 7)\}$.




Figure 3: Vietoris-Rips complexes

¹A more rigorous approach can be found in *Algebraic Topology* citeHatcher




Figure 4: Barcode example

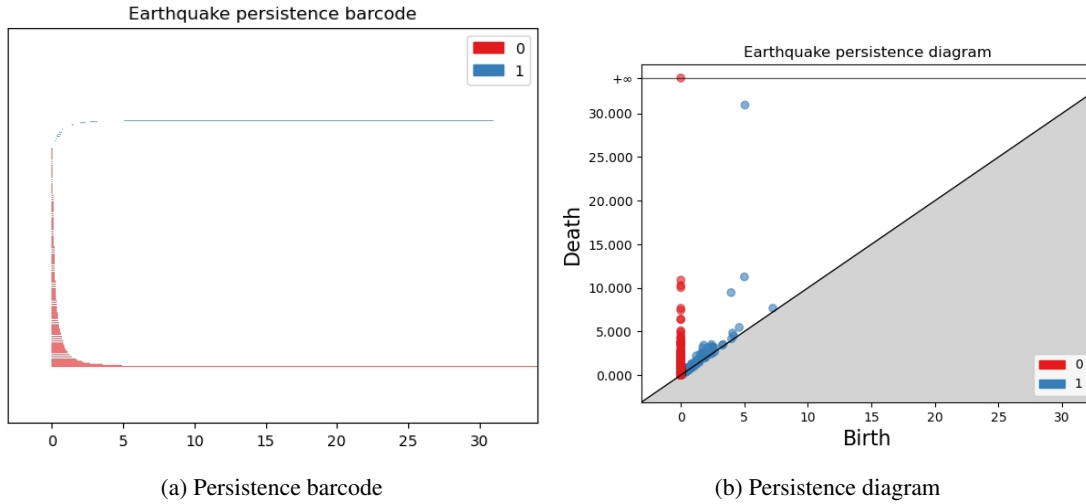


Figure 5: Persistence results on earthquake data

Now, considering all values for $\alpha \geq 0$, we get an evolution of Vietoris-Rips complexes and Betti numbers. Let's start with b_0 . For $\alpha \in [0, 5)$, we have four vertices and no edges, so $b_0 = 4$. However, when $\alpha \in [5, 7)$, we add the edge shown in Figure 3a so two connected components merge into one, and we get $b_0 = 3$. As α increases, points farther from each other get connected by edges and b_0 decreases. We keep track of the changes of b_0 in a **barcode** (Figure 4). Each interval $[b, d)$ represents a feature. b , the birth, is when a feature appears and d , the death, is when it disappears. The barcode is the collection of all these intervals. We observe the interval $[0, 5)$ in the barcode that represents the connected component that lives until we add the first edge.

Similarly, for b_1 , we have a loop that appears at $\alpha = \sqrt{148}$ when we add the edge $[(0, 5), (12, 7)]$. It disappears at $\alpha = 13$ when it becomes the boundary of the two filled triangles. Hence, we get a single interval $[\sqrt{148}, 13)$ for b_1 in Figure 4.

However, barcodes can be difficult to read when studying more complex data. The barcode for the earthquake example is shown in Figure 5a, with the intervals for b_0 in red, and those for b_1 in blue.

The barcode has 1957 intervals for b_0 since we had 1957 points, but many die rapidly, and it is difficult

to differentiate significant features from noise (Figure 5a). An alternate representation is the **persistence diagram**. Each interval $[b, d)$ of the barcode is represented as an ordered pair (b, d) in the plane (Figure 5b). The persistence diagram is the multiset, a set that allows repeated values, of the points on the diagonal, each appearing infinitely many times, as well as those ordered pairs.

There are no points below the diagonal since that would mean the death of a feature happens before its birth. Points close to the diagonal have a short lifespan and are considered noise, so the significant features are the ones with points far from the diagonal. In this diagram, we can easily identify the three primary loops as the blue points far from the diagonal. The blue point that is farther than the others corresponds to the largest loop in the earthquake data since it lives longer.

4 PERSISTENCE LANDSCAPES

Persistence landscapes, introduced by Bubenik [5], are another way of summarizing the topological features of the point cloud. Each point $p = (b, d)$ in the persistence diagram \mathcal{P} is mapped to $(\frac{b+d}{2}, \frac{d-b}{2})$ in the landscape. This transformation corresponds to a clockwise rotation of 45° of the diagram. We then add piece-

wise linear functions of slope ± 1 to each point:

$$\Lambda_p(t) = \begin{cases} t - b & t \in [b, \frac{b+d}{2}] \\ d - t & t \in (\frac{b+d}{2}, d] \\ 0 & \text{otherwise} \end{cases}$$

Then, the landscapes are : $\{\lambda_{\mathcal{P}}(k, t) = k \max_{p \in \mathcal{P}} \Lambda_p(t) \mid t \in [0, T], k \in \mathbb{N}\}$ where the k max is the k^{th} largest value for a specific value of t . The first landscape $\lambda_{\mathcal{P}}(1, t)$ is the regular maximum of functions. Then, after removing the first landscape, we get the second landscape by taking the regular maximum again. We proceed similarly for subsequent landscapes. The peaks in the landscape correspond to significant topological features. Indeed, those would come from points far from the diagonal in the persistence diagram. Figure 6, from “An Introduction to Topological Data Analysis: Fundamental and Practical Aspects for Data Scientists” by Chazal and Michel [1], illustrates the transformation from the diagram to the landscape. The blue curve is the first landscape, the red curve is the second, and the orange is the third. The landscapes for $k \geq 4$ are $\lambda(k, t) = 0(t)$.




Figure 6: Persistence diagram to landscape [1]

The landscape (Figure 7) corresponding to the earthquake data shows the three significant loops as the pronounced blue, orange and green peaks. Since we have 223 points corresponding to 1-dimensional features in the persistence diagram, we will have 223 non-zero landscapes. We only show the 4th and 5th landscapes in red and purple for reference.

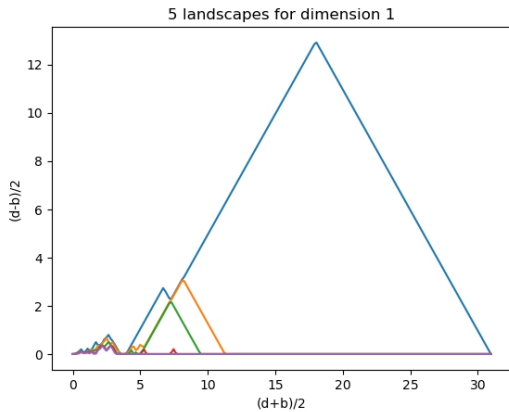


Figure 7: 5 landscapes of dimension 1 from the earthquake data


5 MEAN LANDSCAPE

If we have a large sample, TDA computations can become increasingly slow. A solution is to split it into multiple smaller samples, but we would still like to know the topology of the original sample. Similarly, if we have multiple point clouds sampled from the same population, we aim to gain insight into the topology of the population. A question arises: how should we combine the TDA outputs of the individual samples to understand the whole point cloud better? Currently, there is no definition for an average barcode or an average persistence diagram. However, persistence landscapes have the advantage that many statistical methods can be used on them, and we do have the notion of a mean landscape.

We will denote the first landscape ($k = 1$) obtained from the sample X_i by $\lambda_i(t) = \lambda_i(1, t)$. If we only have one sample X , we can get the X_i 's by resampling X with replacement. This means we give every point in X an equal probability to be chosen, *i.e.* in a sample of size n , every point has $1/n$ chance to be picked. Then, we choose points from X and each selected point is put back into X before choosing the next point.

The first landscape corresponds to the most significant feature throughout the evolution of the Vietoris-Rips complexes with respect to α . The **mean landscape** $\mu(t) = \mathbb{E}(\lambda_i(t))$ is the expected function for the first landscape. It comes from the population space, so we do not know it since we only have access to a sample. We use our samples and their landscapes $\lambda_1, \dots, \lambda_n$ to estimate it by the **average landscape** $\bar{\lambda}_n = \frac{1}{n} \sum_{i=1}^n \lambda_i$. Bubenik showed the pointwise convergence of $\bar{\lambda}_n$ to $\mu(t)$ and the Central Limit Theorem for landscapes [5].

From our sample of 1957 earthquake epicentres, we created 100 samples of size 800 and obtained $\lambda_1, \dots, \lambda_{100}$ (Figure 8a). Figure 8b shows our estimate $\bar{\lambda}_n$ in red. The peak at (19.63, 10.98) corresponds to the most significant 1-dimensional feature: the largest loop. We find that, on average, this loop is born at 8.65 and dies at 30.61.



(a) 1st landscape of the 100 subsamples (b) Average landscape with a 95% confidence band

Figure 8: Landscape results on earthquake data

Beyond finding an estimate for $\mu(t)$, we would like to know its accuracy. Chazal et al. presented an algorithm for finding a uniform band around $\bar{\lambda}_n$, called the **multiplier bootstrap** [2]. This band tells us where our average landscape would most likely be if we repeat our experiment and obtain new samples, depending on our confidence level. For example, a 95% confidence level indicates that approximately 95% of our resampled average landscapes will lie within the band. The idea behind this method is to use the weighted difference between each resampled landscape and the original one to calculate the width of the band.

In *On the Bootstrap for Persistence Diagrams and Landscapes*, Chazal et al. showed the weak convergence of landscapes [4]. We will use this method on landscapes with $k = 1$, but it can be generalized to landscapes with other values of k , by taking the appropriate landscapes as input. Let $\lambda_1, \dots, \lambda_n$ be landscapes, $1 - \alpha$ be the confidence level, and B be the number of bootstrap samples. We follow the steps:

1. Compute $\bar{\lambda}_n(t)$
2. Repeat B times:
 - (a) Generate $\xi_1, \dots, \xi_n \sim N(0, 1)$ from a standard normal distribution
 - (b) $\tilde{\theta}_i = \sup_t \frac{1}{\sqrt{n}} \left| \sum_{i=1}^n \xi_i \cdot (\lambda_i(t) - \bar{\lambda}_n(t)) \right|$
3. $\tilde{Z}(\alpha) = \inf\{z \mid \frac{1}{B} \sum_{i=1}^B I(\tilde{\theta}_i > z) \leq \alpha\}$

Then, the confidence functions are $l_n(t)$ (lower bound) and $u_n(t)$ (upper bound):

$$l_n(t) = \bar{\lambda}_n(t) - \frac{\tilde{Z}(\alpha)}{\sqrt{n}} \quad u_n(t) = \bar{\lambda}_n(t) + \frac{\tilde{Z}(\alpha)}{\sqrt{n}}$$

We applied this method to the earthquake data and our 100 subsamples, with $1 - \alpha = 0.95$ and $B = 100$. We obtain a band of size 0.488 in red, bounded by the confidence functions in black, around the average landscape shown in Figure 8b. The average landscape and the band together give us an estimate of the mean landscape within a margin of error.

6 CONCLUSION

There are multiple summary representations of the topological features of a point cloud such as persistence landscapes that allow us to use statistical methods like the multiplier bootstrap. Using smaller samples makes computations faster, and we can use them

to find an estimate for the mean landscape and build confidence intervals. Thus, the average landscape summarizes some topological features of the population space by combining results obtained from multiple samples.

7 ACKNOWLEDGEMENTS

I want to thank Professor François Charette for his supervision and support. I am very grateful to him for introducing me to topology. I thank the Fonds de recherche du Québec – Nature et technologies (FRQNT) for funding this summer research project (#356843).

REFERENCES

- [1] Chazal, F. & Michel, B. An Introduction to Topological Data Analysis: Fundamental and Practical Aspects for Data Scientists. *Frontiers In Artificial Intelligence*. (2021)
- [2] Chazal, F., Fasy, B., Lecci, F., Rinaldo, A. & Wasserman, L. Stochastic Convergence of Persistence Landscapes and Silhouettes. (2013), <https://arxiv.org/abs/1312.0308>
- [3] Polterovich, L., Rosen, D., Samvelyan, K. & Zhang, J. Topological Persistence in Geometry and Analysis. (2021), <https://arxiv.org/abs/1904.04044>
- [4] Chazal, F., Fasy, B., Lecci, F., Rinaldo, A., Singh, A. & Wasserman, L. On the Bootstrap for Persistence Diagrams and Landscapes. (2014)
- [5] Bubenik, P. Statistical topological data analysis using persistence landscapes. (2015), <https://arxiv.org/abs/1207.6437>
- [6] Project, T. GUDHI User and Reference Manual. (GUDHI Editorial Board, 2024), <https://gudhi.inria.fr/doc/3.10.1/>
- [7] Hatcher, A. Algebraic Topology. (Cambridge University Press, 2002), <https://pi.math.cornell.edu/~hatcher/AT/ATpage.html>
- [8] Kovacev-Nikolic, V., Bubenik, P., Nikolić, D. & Heo, G. Using persistent homology and dynamical distances to analyze protein binding. *Statistical Applications In Genetics And Molecular Biology*. **15** (2016)

STATISTICAL INFERENCE FOR SURVIVAL DATA

Jiajun Zhang

Survival analysis studies the behavior of time-to-event data, which arise naturally in modern statistical studies and are characterized by incomplete observations due to censoring and truncation. These unique features invalidate classical statistical methods and require the development of specialized inferential tools. This article presents an introduction to basic theories and methodologies used for modeling time-to-event data, which are widely used in biostatistical and epidemiological research. We begin by introducing the censoring and truncation mechanism, followed by defining basic quantities, including survival and hazard functions. We next study the key estimating techniques for survival data based on parametric, non-parametric, and semi-parametric inferences. We briefly introduce likelihood inference for parametric inference, life table methods, the Kaplan-Meier (KM) estimator and its variants for non-parametric inference, and Cox's proportional hazard (PH) model for semi-parametric inference. Finally, we introduce the log-rank test, which is commonly used for testing statistical significance in the survival data among multiple groups.

1 INTRODUCTION

Time-to-event data often presents characteristic features known as censoring and truncation. Censoring occurs when the observation is incompletely determined for some subjects, and truncation occurs when only a group of chosen individuals are observed and the rest are systematically excluded from the study.

There are three main types of censoring: Right censoring, left censoring, and interval censoring. Right censoring happens when we only know that the event happens after a certain time threshold; Left censoring happens when we only know that the event time happens before a certain time threshold; Interval censoring occurs when we only know that the event time happens during a certain time period. Below, we will use three examples taken from Klein and Moeschberger [2] to illustrate different types of censoring.

Example. During 1982-1992, 863 patients underwent kidney transplantation at the Ohio state university transplant center, where researchers investigated the survival time of those patients after transplantation. Their maximum follow-up time was 9.47 years (Right censoring).

Example. In 1978, Turnbull and Weiss studied the first usage of marijuana among California high school boys, and some answered with "I used it before, but I cannot remember when" (Left censoring).

Example. The National Longitudinal Survey of Youth investigated a random sample of females aged 14 to 21 yearly from 1979 to 1988. Females in the survey were asked about any pregnancies that have occurred since they were last interviewed (Interval censoring).

Truncation, on the other hand, can be viewed as a design issue where a group of individuals from the target population is systematically excluded from the

sampling pool. This feature becomes obvious in survival analysis because, due to cost, we do not have access to all the individuals. For those individuals who died before entering the study, we do not have any access to their exact death time, we say those individuals were left truncated.

2 BASIC QUANTITIES OF SURVIVAL DATA

2.1 Survival Function and Hazard Function

Suppose we have drawn a random sample (that is, the life time of those individuals are mutually independent and identically distributed) of n individuals. We will use X_i as the random variable denoting the life time of the i th individual. We use $f(x)$ as the probability density function (pdf) of X_i , $F(x)$ as the cumulative distribution function (cdf) of X_i . We first define the **survival function** $S(x)$:

Definition. Let X be a non-negative random variable which denotes the life time, let $F(x)$ be the cumulative distribution function (cdf) of X , then the survival function is defined as

$$S(x) = \mathbb{P}(X \geq x) = 1 - F(x) = 1 - \int_0^x f(t)dt.$$

The survival function gives the probability that a subject will survive beyond the indicated time $t = x$. By the uniqueness of a distribution function, the survival function also uniquely defines a distribution. Another characteristic is called the hazard function $h(x)$, which is the instantaneous rate of death at time right after $t = x$ given that the individual is still alive at $t = x$.

Definition. Let X be a non-negative random variable that denotes the life time, then the hazard function $h(x)$ is defined as

$$h(x) = \lim_{dx \rightarrow 0} \frac{\mathbb{P}(x \leq X \leq x + dx | X \geq x)}{dx}.$$

From Definition 2.1, we can further derive the hazard function as the ratio of the probability density function (pdf) and the survival function:

$$\begin{aligned} h(x) &= \lim_{dx \rightarrow 0} \frac{\mathbb{P}((x \leq X \leq x+dx) \cap (X \geq x))}{\mathbb{P}(X \geq x)dx} \\ &= \lim_{dx \rightarrow 0} \frac{\mathbb{P}(x \leq X \leq x+dx)}{[1-F(x)]dx} \\ &= \lim_{dx \rightarrow 0} \frac{F(x+dx) - F(x)}{[1-F(x)]dx} \\ &= \frac{F'(x)}{1-F(x)} = \frac{f(x)}{S(x)}. \end{aligned}$$

In addition, if we apply the Chain rule, we have

$$h(x) = -\frac{d}{dx} \log(S(x)),$$

by integrating both sides, we have

$$H(x) = \int h(x)dx = -\log S(x), \quad (1)$$

where the integral term is denoted as the **cumulative hazard function**.

2.2 Parametric Inferences with Censoring and Truncation

After we have selected our random sample, the goal is to find the correct probability distribution for the survival function $S(x)$ or the hazard function $h(x)$, so we are able to make further inferences on it. In statistics, we have three different types of inferences on the random sample: Parametric, non-parametric, and semi-parametric. Their characteristics are as follows.

- **Parametric Inference:** The distribution of the sample is unknown up to finitely many parameters. For example, we know X forms an exponential distribution, but with parameter θ unknown. The methodologies are the easiest but less common in practice.
- **Non-parametric Inference:** We do not have any information on the distribution of the sample, and is very common in practice.
- **Semi-Parametric Inference:** This combines the parametric and non-parametric inference, where the distribution of the sample consists of a parametric part and a non-parametric part.

In this section we will focus on parametric inference. We will discuss non-parametric and semi-parametric inferences in section 3 and 4.

Without any censoring or truncation, our task would be a straightforward parametric estimation

problem where maximum likelihood estimation (MLE) method is widely used. With censoring and truncation introduced, we shall take them into consideration, and the goal of this section is to construct the corresponding inferences. We first introduce the likelihood function:

Definition. Suppose we have a random sample of $X_1, \dots, X_n \stackrel{i.i.d}{\sim} f(x, \theta)$ where f is the probability density function (pdf) of the sample, θ is the unknown parameter of interest. The **likelihood function** $\mathcal{L}(\theta)$ is given by

$$\mathcal{L}(\theta) = \prod_{i=1}^n f(x_i, \theta).$$

Once we have obtained the likelihood function, the **maximum likelihood estimate** $\hat{\theta}$ of θ is obtained by maximizing $\mathcal{L}(\theta)$ as a function of θ . We first consider the presence of (right) censoring: Suppose for each individual in the sample, X_i is the true life time, C_i is the censoring time independent of X_i . We observe (T_i, δ_i) , where $T_i = \min\{X_i, C_i\}$, and

$$\delta_i = \begin{cases} 1 & \text{if } X_i \leq C_i \\ 0 & \text{if } X_i > C_i \end{cases}. \quad (2)$$

If $\delta_i = 1$, it means the life time is less than the censoring time, and hence we are able to observe the exact life time, and we may directly use $f(x_i, \theta)$ as the likelihood for X_i . On the other hand, if $\delta_i = 0$, it means the individual is (right) censored. In this case we know $X_i \geq C_i$ hence the likelihood is proportional to $\mathbb{P}(X_i \geq C_i)$, given by $S(C_i, \theta)$. The total likelihood using equation (2) is given by

$$\mathcal{L}(\theta) \propto \prod_{i=1}^n [f(x_i, \theta)]^{\delta_i} \cdot [S(C_i, \theta)]^{1-\delta_i}.$$

As for truncation, what we observe becomes the conditional probability and conditional distribution. We usually work with left truncation, meaning that we are only able to observe the samples which satisfy $X_i > H$ for the pre-determined truncation time H . In this case, the distribution of X_i is the conditional distribution of X_i given $X_i > T$, i.e.,

$$f(X_i = x_i | X_i > H) = \frac{f(x_i, \theta)}{S(H, \theta)},$$

thus the likelihood is given by

$$\mathcal{L}(\theta) = \prod_{i=1}^n \frac{f(x_i, \theta)}{S(H, \theta)}.$$

3 LIFE TABLE AND KAPLAN MEIER ESTIMATOR

3.1 Non-Parametric Inference

In reality we often have non-parametric inferences where we do not know the distribution of the sample in advance: All we have is just a collection of observable data. We will try to make inference and estimate key survival features based on non-parametric inferences.

Definition. Given a random sample X_1, \dots, X_n , we define the **empirical cumulative distribution function (ECDF)** F_n by

$$F_n(x) = \frac{1}{n} \sum_{i=1}^n \mathbf{1}\{X_i \leq x\}, \quad (3)$$

where

$$\mathbf{1}\{X_i \leq x\} = \begin{cases} 1 & \text{if } X_i \leq x \\ 0 & \text{otherwise} \end{cases}.$$

$F_n(x)$ estimates the true distribution $F(x)$ by the ratio of the individuals in our observable sample that satisfies $X_i \leq x$. Based on this, the survival function $S(x)$ can then be estimated by

$$\hat{S}(x) = \frac{1}{n} \sum_{i=1}^n \mathbf{1}\{X_i \geq x\}. \quad (4)$$

Using **Glivenko-Cantelli theorem**, the ECDF and empirical survival function converges to their population counterparts. However, this is based on the assumption that no censoring or truncation is present. In our setting, equations (3) and (4) are not good estimates. In the rest of this section, we discuss two other popular non-parametric methods: Life table and Kaplan-Meier estimator.

3.2 Life Table

Life table is one of the earliest methods to estimate the survival distribution non-parametrically. It is a table which shows the survival data of a certain population at a certain age (or a certain time period). It represents the survivorship of people in a population, or the population's longevity. It usually has several components, where each column will represent a certain survival feature. See table 2 in the appendix for an example.

There are two types of life tables, namely **current life tables** and **cohort life tables**. A current life table measures the mortality of the population at a specific time period, while a cohort life table measures the mortality of the population at a specific age. That is, in a cohort life table, all individuals from the population are

believed to have the same date (period) of birth, like table 2 in the appendix. While in a current life table, the ages of the individuals may vary, and it would not be a great representation of the mortality behavior of the individuals at a certain age. Hence throughout the discussion, we will focus on cohort life table and we will see how we can estimate key survival quantities from it.

As an example, table 2 in the appendix shows the cohort life table of Canadian populations between age 60 and 80 from year 2021 to 2023. We define l_x as the number of survivors at age x (in Table 2, l_{60} would be 91,805). So the number of deaths in year x would be $d_x = l_x - l_{x+1}$. If we start our study from year 0, then the survival probability at age x can be estimated by

$$\begin{aligned} \hat{S}_x &= \frac{l_x}{l_0} \\ &= \frac{l_1}{l_0} \times \frac{l_2}{l_1} \times \dots \times \frac{l_x}{l_{x-1}} \\ &= \prod_{i=0}^x \left(1 - \frac{d_i}{l_i}\right). \end{aligned} \quad (5)$$

3.3 Kaplan-Meier Estimator

We introduce an estimator Kaplan and Meier proposed, known as the **Kaplan-Meier (KM) Estimator**.

Theorem 1 (Kaplan-Meier [1], 1958). Suppose we have observed the ordered event time t_1, \dots, t_m , where d_j deaths are observed at t_j , n_j individuals are at risk (event still not occurred) at t_j^- , then the Kaplan-Meier (KM) estimator of the survival function $S(t)$ is given by

$$\hat{S}(t) = \prod_{t_j \leq t} \left(1 - \frac{d_j}{n_j}\right). \quad (6)$$

We propose a plug in approach to derive the KM estimator given in equation (6), where we split the time into sub-intervals: $[0, t_1), [t_1, t_2), \dots, [t_m, +\infty)$. If $t \in [0, t_1)$, it means that no deaths had been observed in that interval, and in this case the survival probability is simply 1. Next if $t \in [t_1, t_2)$, by definition, we have

$$\begin{aligned} S(t) &= \mathbb{P}(X \geq t_1) \\ &= \mathbb{P}(\text{Survived after } t_1) \\ &= \frac{\text{Number of Individuals survived after } t_1}{\text{Number of individuals at risk at } t_1^-} \\ &= \frac{n_1 - d_1}{n_1}. \end{aligned}$$

Now we consider $t \in [t_2, t_3)$, then we have the condi-

tional probability given by

$$\begin{aligned} S(t) &= \mathbb{P}(X \geq t_2) \\ &= \mathbb{P}(\text{Survived after } t_2) \\ &= \mathbb{P}(\text{Survived after } t_2 \mid \text{Survived after } t_1) \cdot \mathbb{P}(\text{Survived after } t_1) \\ &= \frac{\text{Number of Individuals survived after } t_2}{\text{Number of individuals at risk at } t = t_2^-} \times \frac{n_1 - d_1}{n_1} \\ &= \frac{n_2 - d_2}{n_2} \times \frac{n_1 - d_1}{n_1}. \end{aligned}$$

Thus, we may use this plug in approach recursively, and we will get the KM estimator:

$$\hat{S}(t) = \prod_{t_j \leq t} \left(\frac{n_j - d_j}{n_j} \right) = \prod_{t_j \leq t} \left(1 - \frac{d_j}{n_j} \right),$$

which is the same as equation (6). Since our survival function is based on the terms of products, this is also known as the product limit (PL) estimator. As we can see, Kaplan-Meier estimator is very similar to equation (5) obtained by using life table, however, the key difference between them is that, in KM estimator, we have observed the exact time of occurrence t , while in the life table method, the sample is interval censored.

Kaplan-Meier estimator is **almost unbiased**. The unbiasedness holds when assuming n_j is always non-zero. But despite that, the true bias is exponentially small [3]. The variance of the Kaplan-Meier estimator can be computed by **Greenwood's formula**:

Theorem 2 (Greenwood's formula). *Let $\hat{S}(t)$ be the KM estimator, then*

$$\widehat{\text{Var}}(\hat{S}(t)) = \hat{S}^2(t) \cdot \sum_{t_j \leq t} \frac{d_j}{n_j(n_j - d_j)},$$

where d_j is the number of death observed at t_j , n_j is the number of individuals at risk at t_j^- , as in the KM estimator.

Assuming $n_j \neq 0$, we are able to derive the $100(1 - \alpha)\%$ confidence interval by **Central limit theorem**:

$$\hat{S}(t) \pm z_{\alpha/2} \sqrt{\widehat{\text{Var}}[\hat{S}(t)]}. \quad (7)$$

A log-log transformation [9] can be used to adjust the upper and lower bound of equation (7) to within $[0, 1]$.

3.4 Variants of KM Estimator

There are several variants of the KM estimator, including **Nelson-Aalen (NA) estimator** and **Aalen-Johansen (AJ) estimator**. NA estimator estimates the cumulative hazard function $H(x)$ defined through equation 1, where we have

$$H(t) = -\log(S(t)),$$

by replacing $S(x)$ with KM estimator, we have

$$\hat{H}(t) = -\sum_{t_j \leq t} \log \left(1 - \frac{d_j}{n_j} \right) = \sum_{t_j \leq t} \frac{d_j}{n_j} + \mathcal{O} \left(\frac{d_j^2}{n_j^2} \right).$$

When the number of events at each time is small relative to the risk set ($d_j \ll n_j$), we may simply drop the higher order term, and hence the cumulative hazard function can be estimated by

$$\hat{H}(t) = \sum_{t_j \leq t} \frac{d_j}{n_j}. \quad (8)$$

where equation (8) is known as the NA estimator.

AJ estimator can be viewed as a generalized version of KM estimator where the model has several intermediate states rather than simply alive and death. A common model shown in figure 1 is a healthy-diseased-death model:




Figure 1: A multi-state model

We denote 0 as the “healthy state”; 1 as the “diseased state”, and 2 as the “death state”. In figure 1, there are a total of three possible transitions, and we denote $\alpha_{ij}(t)$ as the instantaneous risk of experiencing a transition from state i to j at time t . This can be viewed as a generalized hazard function between two states. We also define $\mathbb{P}_{ij}(s, t)$ as the probability that an individual in state i at time s will be in state j at time t , which can also be viewed as a generalized version of survival function between two states. For example, the probability of staying healthy from time s to t can be theoretically computed by

$$\mathbb{P}_{01}(s, t) = \exp \left\{ -\int_s^t (\alpha_{01}(x) + \alpha_{02}(x)) dx \right\}.$$

In practice, we replace α_{01}, α_{02} with their corresponding NA estimator:

$$\hat{\mathbb{P}}_{01}(s, t) = \prod_{s < t_j < t} \left(1 - \frac{d_{0j}}{n_{0j}} \right),$$

where d_{0j} represents the total number of transitions from state 0 at time t_j , and n_{0j} is the total number of individuals at risk at time t_j^- .

We may use similar methods to provide estimates for $\mathbb{P}_{11}(s, t), \mathbb{P}_{12}(s, t)$, etc., but it is wise to move one

step further and consider a more general case where $\mathcal{I} := \{0, 1, \dots, k\}$ denotes $k + 1$ different states. Define $\alpha_{gh}(t)$ as the instantaneous risk of transition from state g to state h ($g \neq h$) at time t . Define $\mathbb{P}_{gh}(s, t)$ to be the probability that an individual at state g at time s will be at state h at time t , and we define the **transition matrix** as

$$\mathbf{P}(s, t) = \begin{bmatrix} \mathbb{P}_{00} & \mathbb{P}_{10} & \cdots & \mathbb{P}_{k0} \\ \mathbb{P}_{01} & \mathbb{P}_{11} & \cdots & \mathbb{P}_{k1} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbb{P}_{0k} & \mathbb{P}_{1k} & \cdots & \mathbb{P}_{kk} \end{bmatrix}.$$

We denote d_{ghj} as the number of individuals with transition from state g to state h at observed time t_j , and $d_{gj} = \sum_{h \neq g} d_{ghj}$ as the total number of transitions out of state g at observed time t , n_{gj} as the number of individuals at state g at time t_j^- . Finally, we define the $(k + 1) \times (k + 1)$ matrix $\hat{\alpha}_j$ with entries (g, h) by

$$\hat{\alpha}_j(g, h) = \begin{cases} \frac{d_{ghj}}{n_{gj}} & \text{if } g \neq h \\ -\frac{d_{gj}}{n_{gj}} & \text{if } g = h \end{cases},$$

then the AJ estimator [4] takes the form

$$\hat{P}(s, t) = \prod_{s < t_j < t} (\mathbf{I} + \hat{\alpha}_j),$$

where \mathbf{I} is the $(k + 1) \times (k + 1)$ identity matrix.

3.5 Counting Processes

The properties of the estimators we introduced can be formally studied using counting processes and martingale theory. We will denote the number of death, $D(t)$, as a **counting process** such that $D(0) = 0, D(t) - D(s)$ is the number of deaths in the interval $[s, t)$. If we denote $h(t)$ by the true hazard rate, $N(t)$ by the number of individuals at risk at t , then the predicted death at t is $N(t) \cdot h(t)$, and the total predicted death up to t is given by $\int_0^t N(s) \cdot h(s) ds$. It follows from **Doob-Meier decomposition** that the difference between total observed death and predicted death is a mean zero martingale:

$$M(t) = D(t) - \int_0^t N(s) \cdot h(s) ds \quad \text{where } \mathbb{E}[M(t)] = 0. \tag{9}$$

It follows from equation (9) that

$$\frac{dM(t)}{N(t)} = \frac{dD(t)}{N(t)} - h(t),$$

while integrating again, we have

$$H(t) := \int_0^t h(t) dt = \int_0^t \frac{dD(t)}{N(t)} - \int_0^t \frac{dM(t)}{N(t)}, \tag{10}$$

The Nelson-Aalen estimator can also be defined using a Riemann-Stieltjes integral: $\hat{H}(t) = \int_0^t \frac{dD(t)}{N(t)}$, which corresponds the one in equation (10), hence we see the unbiasedness of NA estimator:

$$\mathbb{E}[\hat{H}(t) - H(t)] = \mathbb{E} \left[\int_0^t \frac{dM(t)}{N(t)} \right] = 0,$$

since $M(t)$ is a mean zero martingale. The unbiasedness of KM estimator by a similar way. A rigorous approach can be found in Fleming and Harrington [7], and Li [8].

4 COX'S PROPORTIONAL HAZARD (PH) MODEL

4.1 Model Set Up

Cox's proportional Hazard Model [5] is a semi-parametric model which utilizes **covariates**, for example the individual's gender, age, habits, etc. We usually denote $\vec{z} \in \mathbb{R}^p$ as a p -dimensional covariate vector, and the survival probability becomes the conditional probability given the covariates \vec{z} . Cox's PH model suggests that the hazard function given the covariates \vec{z} takes the form

$$h(x|\vec{z}) = h_0(x) \cdot \exp(\vec{z}^\top \vec{\beta}), \tag{11}$$

where $h_0(x)$ is the **baseline hazard**, and $\vec{\beta}$ is the **coefficient vector**. To estimate the coefficients vector $\vec{\beta}$, we assume for simplicity that there are no tied events (which means at each t_j there is exactly one death observed). Suppose at time t_0 we have m individuals with hazard function $h_1(t_0), \dots, h_m(t_0)$ respectively, and we know there is exactly one death, then the probability that individual j has the event is given by

$$\mathbb{P}(\text{individual } j \text{ has the event} | t_0) = \frac{h_j(t_0)}{h_1(t_0) + \dots + h_m(t_0)}.$$

If we replace each $h_i(t_0)$ by equation (11), we have

$$\mathbb{P}(\text{individual } j \text{ has the event} | t_0) = \frac{\exp(\vec{z}_j^\top \vec{\beta})}{\sum_{i=1}^m \exp(\vec{z}_i^\top \vec{\beta})},$$

where the baseline hazard is not required. We see that this probability is independent of the event time, and only depends on the order in which events occur. Based on this observation we specify a model where we assume individual i fails at time t_i exactly. Then by multiplying the probability of each individual, we get the **partial likelihood** as follows:

$$\mathcal{L}(\vec{\beta}) = \prod_{i=1}^m \frac{\exp(\vec{z}_i^\top \vec{\beta})}{\sum_{j \in R(t_i)} \exp(\vec{z}_j^\top \vec{\beta})}, \tag{12}$$

where $R(t_i)$ is the set of individuals at risk at time t_i . Likelihood inference can be based on equation (12), and one may estimate the coefficient vector β by maximizing the function given by equation (12).

$$\mathcal{L}(\hat{\beta}) = \sup_{\beta} \prod_{i=1}^m \frac{\exp(\bar{z}_i^\top \beta)}{\sum_{j \in R(t_i)} \exp(\bar{z}_j^\top \beta)}. \quad (13)$$

4.2 Relative Hazard

One advantage of Cox’s PH model is that the exponential part is independent of time, and thus remains a constant under the ratio. If we choose two individuals with different covariates \bar{z}_1, \bar{z}_2 , the ratio of the two hazard rates is a constant given by

$$\frac{h(x|\bar{z}_1)}{h(x|\bar{z}_2)} = \exp\left((\bar{z}_1 - \bar{z}_2)^\top \beta\right).$$

Hence despite not knowing the baseline hazard $h_0(x)$, we may compare the ratio of different covariates, and see the relative hazard rate among reference group and non-reference group.

We next present a study of 863 kidney transplant patients [2] where two different covariates were introduced: The patients’ gender (male and female) and race (white and black). Then we have 4 groups, namely black male; black female; white male; white female. We set $z_1 = 1$ if we have a black male and 0 otherwise; $z_2 = 1$ if we have a white male and 0 otherwise; $z_3 = 1$ if we have a black female and 0 otherwise. The white female is considered as the reference group and its covariate vector is simply $\bar{z} = \vec{0}$, and its hazard function in Cox model is equal to the base line hazard $h_0(t)$, and we have

$$h(t|\bar{z}) = h_0(t) \exp\left(\beta_1 z_1 + \beta_2 z_2 + \beta_3 z_3\right).$$

By solving equation (13) numerically, we have $\beta_1 = 0.1596, \beta_2 = 0.2484, \beta_3 = 0.6567$. Now we can see the performances of non-reference group compared to the reference group:

$$\frac{h(t|z_1)}{h_0(t)} = 1.17, \frac{h(t|z_2)}{h_0(t)} = 1.28, \frac{h(t|z_3)}{h_0(t)} = 1.93.$$

Thus the relative hazard risks for black male, white male and black female compared to white female are given by 1.17, 1.28, 1.93 respectively. One can also use this method to test if a certain treatment is effective by defining the group with placebo as the reference group, and see the relative hazard risk for treatment group. If the hazard rate is greater than 1, it means the treatment is somehow not effective, if less than 1 it means the treatment is somehow effective.

5 HYPOTHESIS TESTING

In this section, we will introduce the **log-rank test**, which is a non-parametric hypothesis testing methods commonly used. It can be used to compare the survival data from 2 or more groups (most commonly, treatment group & placebo group), and conclude their statistical similarity. This usually involves a null hypothesis \mathcal{H}_0 : They are statistically equivalent at a certain significance level, and an alternative hypothesis \mathcal{H}_1 : They are not statistically equivalent at a certain significance level.

We will consider a study done by Freireich et. al. [6], where they studied the effectiveness of 6-mercaptopurine (6-MP) among children with leukemia. In their experiment, 21 children were given 6-MP treatment (group 1), and 21 children were given a placebo (group 2). The data collected is shown in Table 3 in the appendix.

We use KM and NA estimators to estimate the survivor and hazard function for these data and generate the following plots:




Figure 2: Kaplan–Meier survival curve with 95% confidence interval.

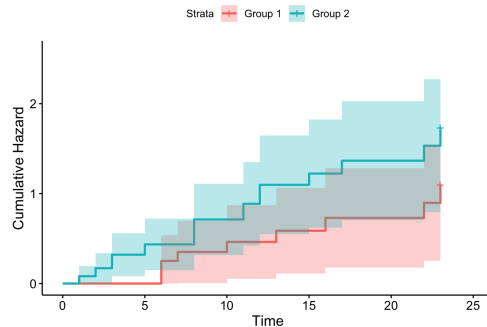


Figure 3: Nelson-Aalen hazard curve with 95% confidence interval.

From the plots we can see that group 1 (6-MP) has a lower hazard rate and higher survival probability

compared to group 2 (placebo), and hence we may conclude that 6-MP is important in leukemia remission. In order to conclude our hypothesis, a rigorous testing is required. We first define the expected number of failures e_{if} for group $i = 1, 2$ at each observed failure time by

$$e_{1f} = \left(\frac{n_{1f}}{n_{1f} + n_{2f}} \right) \times (d_{1f} + d_{2f}),$$

$$e_{2f} = \left(\frac{n_{2f}}{n_{1f} + n_{2f}} \right) \times (d_{1f} + d_{2f}).$$

That is, e_{if} is the proportion of individuals at risk in group i multiplied by the total number of failures over both groups. Let $n_f = n_{1f} + n_{2f}$ and $d_f = d_{1f} + d_{2f}$, at time $t_{(f)}$, we obtain the following table:

	Group 1	Group 2	Total
Fails	d_{1f}	d_{2f}	d_f
Survivals	$n_{1f} - d_{1f}$	$n_{2f} - d_{2f}$	$n_f - d_f$
Total	n_{1f}	n_{2f}	n_f

Table 1: contingency table for all subjects in the risk set at time $t_{(f)}$.

Under the null hypothesis \mathcal{H}_0 , where the survival curve of group 1,2 are statistically equivalent, the random variable d_{if} forms a hyper-geometric distribution, with expected value and variance given by

$$\mathbb{E}[d_{if}] = e_{if} = \frac{n_{if}}{n_f} \cdot d_f$$

$$\mathbf{Var}[d_{if}] = v_{if} = n_{if} \cdot \frac{d_f}{n_f} \cdot \frac{n_f - d_f}{n_f} \cdot \frac{n_f - n_{if}}{n_f - 1}.$$

We further define the total difference between observed fails and the expected fails in each group by

$$D_i - E_i = \sum_{f=1}^n (d_{if} - e_{if}), i = 1, 2,$$

then the **log-rank statistic** is defined by

$$Z^2 = \frac{(D_i - E_i)^2}{\mathbf{Var}(D_i)},$$

for any group $i = 1, 2$. Under \mathcal{H}_0 , the log-rank statistic Z^2 forms a χ -squared distribution with one degree of freedom. From table 3 we are able to compute the

log-rank statistic directly: $Z^2 = 10.37$. The p -value is given by 0.0013, which provides strong evidence against the null hypothesis. Hence there is a significant difference in leukemia remission between 6-MP group and placebo group.

6 ACKNOWLEDGMENTS

The original research was supervised by Professor Masoud Asgharian at the department of Mathematics and Statistics at McGill University. This article summarizes my original research report *Analyzing Incident and Prevalent Cohort Survival Data* written in Summer 2025. I would like to thank Professor Masoud Asgharian for his professional guidance, without whose support this article would not have been possible.

REFERENCES

- [1] Kaplan, E. & Meier, P. Non-parametric Estimation from Incomplete Observations . (Journal of the American Statistical Association,1958)
- [2] Klein, J. & Moeschberger, M. Survival Analysis: Techniques for Censored and Truncated Data, 2003. (Springer,2003)
- [3] Zhou, M. Two-Sided Bias Bound of the Kaplan–Meier Estimator . (1988)
- [4] Borgan, O. Aalen–Johansen Estimator. . (2005)
- [5] Cox, D. Regression Models and Life Tables. . (Journal of the Royal Statistical Society. Series B (Methodological),1972)
- [6] Kleinbaum, D. & Klein, M. Survival Analysis, A Self-Learning Text. . (Springer,2012)
- [7] Fleming, T. & Harrington, D. Counting Processes and Survival Analysis. . (Wiley Series in Probability,1991)
- [8] Li, Y. Lecture Notes on Survival Analysis: A Counting Processes Approach. .
- [9] Breheny, P. Survival Data Analysis, BIOS:7210/STAT:7570/IGPI:7210 Lecture Notes. . (2019)

APPENDIX

	A	B	C	D
Age	2021 to 2023	2021 to 2023	2021 to 2023	2021 to 2023
60	91,805	552	0.00601	24.95
61	91,253	598	0.00655	24.10
62	90,655	648	0.00715	23.26
63	90,007	703	0.00781	22.42
64	89,304	763	0.00854	21.59
65	88,541	828	0.00935	20.77
66	87,713	898	0.01024	19.96
67	86,815	975	0.01123	19.17
68	85,840	1,058	0.01233	18.38
69	84,782	1,149	0.01355	17.60
70	83,634	1,246	0.01490	16.84
71	82,387	1,352	0.01641	16.08
72	81,035	1,466	0.01809	15.34
73	79,570	1,588	0.01996	14.62
74	77,982	1,719	0.02204	13.90
75	76,263	1,858	0.02437	13.21
76	74,405	2,007	0.02697	12.52
77	72,398	2,163	0.02988	11.86
78	70,235	2,327	0.03313	11.21
79	67,908	2,498	0.03678	10.57
80	65,410	2,674	0.04087	9.96

$t_{(f)}$	d_{1f}	d_{2f}	n_{1f}	n_{2f}
1	0	2	21	21
2	0	2	21	19
3	0	1	21	17
3	0	2	21	16
5	0	2	21	14
6	3	0	21	12
7	1	0	17	12
8	0	4	16	12
10	1	0	15	8
11	0	2	13	8
12	0	2	12	6
13	1	0	12	4
15	0	1	11	4
16	1	0	11	3
17	0	1	10	3
22	1	1	7	2
23	1	1	6	1

Table 2: (Left) The Life Table of Canadian Population Between Age 60 to 80 from year 2021 to 2023. The elements in the first rows represent: (A) The number of survivors at age x ; (B) The number of deaths between age x and $x + 1$; (C) Death probability between age x and $x + 1$; (D) Life expectancy (in years) at age x . *Source: Statistics Canada. Table 13-10-0114-01 Life expectancy and other elements of the complete life table, three-year estimates, Canada, all provinces except Prince Edward Island.*

Table 3: (Right) The remission data among the sample of 42 children. For each ordered time $t_{(f)}$ in the sample, we use d_{if} to denote the number of children failed (whether leukemia relapse or death) at that time, separately by group i ; we use n_{if} as the number of children at risk at that time. [6]

RECURRENCE ON THE FROG MODEL

Angela Ji and Yanshuo Liu

Our expository report aims to introduce the frog model, which is an infinite system of interacting random walks, taking place on a d -ary tree. The frog model has applications in simulating the spread of disease and information. We would also like to explore some results on the recurrence of the frog model by investigating the question “what is the probability the frog model on a d -ary tree will revisit the origin an infinite amount of times?” We do so by summarizing key results from the paper “Recurrence and Transience for the Frog Model on Trees” by Hoffman, Johnson and Junge, published in the 2017 edition of the *Annals of Probability*. We will show recurrence for $d \leq 2$ by using a step-by-step coupling of the frog model with a simpler version of the model. We then stochastically dominate the model using a functional operator on a set of Poisson variables. We aim to break down the results so that someone with an undergraduate level knowledge of probability and statistics can understand the proofs and how the model works.

1 INTRODUCTION

A **random walk** is a stochastic process that describes the path of a particle. The path consists of a random sequence of steps.

Definition. A **d -ary tree** (denoted T_d) is an infinite rooted graph which starts with a root (origin) and has d branches from each node.

The frog model begins with an awake frog at the root, and a sleeping frog at every node below it. The awake frog moves along the tree in a simple random walk and wakes the frog at every node it stops at. The initially sleeping frog at that node is now awake and goes on its own random walk on the tree, waking all frogs it reaches as well. Frogs can go both forward and backward. Each frog, not starting at the origin, has a $\frac{d-1}{d}$ probability of going down the tree, and a $\frac{1}{d}$ probability of going up the tree.

Remark. The frogs can also be called particles. Branches that come from a node of the tree are called that node’s children.

Remark. The rooted 2-ary (binary) tree (T_2) has a degree-2 root and degree-3 other nodes, enabling recursive frog model analysis through generating functions.

The frog model has applications in modeling the spread of disease and the spread of information. It can be used to investigate threshold effects and network exploration. The frog model is part of a family of self-interacting random walks that has garnered some interest in the research of stochastic processes. Some examples include a general survey on the model done by Robin Pemantle [5], an investigation into its recurrence on Z_d by Telcs and Wormald [6] and exploring the relationship between phase transition and density by Alves et al. [4]

Definition. If the total number of visits to the root by all frogs combined (including the original and all awakened frogs) is infinite, then the model is called **recurrent**. This implies that an infinite number of distinct frogs visit the root. Otherwise, it is called **transient**

The following theorem and all results within its proof are proven in Theorem 1(i) of Recurrence and Transience for the Frog Model on Trees by Hoffman, et al. [3]

Theorem 1. The frog model is recurrent when $d \leq 2$.

There is strong evidence that the frog model is recurrent when $d = 3$ but those results are outside the scope of this report. We will also not be investigating the transience of the frog model, which can be proved for $d \geq 5$, and is likely true for $d = 4$.

Definition. Let A and B be real-valued random variables. We say that A **stochastically dominates** B if

$$\mathbb{P}(A \geq x) \geq \mathbb{P}(B \geq x) \quad \forall x \in \mathbb{R}. \quad (1)$$

We will thus prove Theorem 1 using the notion of stochastic dominance. For the proof of theorem we will couple the frog model with a process where the root is visited less often. If we let V denote the amount




Figure 1: T_2 tree.

of visits to the root in the simpler model, we will see that V dominates a *Poisson* with any mean, and thus goes to ∞ .

Definition. A *branching random walk* is a random walk where each particle produces a number of offspring at every step. The number of offspring produced is often governed by a probability distribution. After branching, all particles and offspring continue on the random walk.

2 RECURRENCE

2.1 Simplified Models

We use a **non-backtracking model** to make recurrence analysis tractable. The model is different from the original one since the frogs move randomly but cannot return to their immediately previous vertex. In this case, the frog model will preserve recurrence and transience while reducing path dependencies.

We will also use the **self-similar model**, which behaves as a sub-process to the original. To achieve this, denote $T_d(v)$ as the subtree of T_d rooted at v , such that we have the non-root vertex v' and its parent v .

By the image of the original tree and the self-similar frog model, we conclude that:

1. Only one frog is allowed to enter the subtree.
2. Each subtree is a smaller copy of the full tree.
3. Arbitrarily choose one frog to continue moving on the tree and stop the remainder of the frogs at v .




Figure 2: Self-similar frog model

Proposition 2. Let T_d be the tree on which the self-similar frog model is defined. For any vertex v , the frog model restricted to the subtree rooted at v has the same distribution as the original frog model on T_d .

Proof. Denote an arbitrary subtree $T_d(v')$ and ignore the frog outside it (which is just the rule: If there are more than one frogs move from v to v' , stop all of them but arbitrary one at v').

By the results above, the self-similar frog model allows only one frog each time to enter the subtree rooted by v' , and all frogs outside the subtree are no longer moving thus preventing external interference. The result is an identical copy of the original model. \square

2.2 Coupling the models

In this section, we have a step-by-step coupling, which is:

Original model \rightarrow Non-backtracking model \rightarrow
Self-similar model

For each original model, construct a non-backtracking model by taking a random walk with deleting the ‘backtracks’, and in this case, there will be fewer return to the root and thus if we can show that the non-backtracking visits (to the root) are finite, then can directly conclude that the original one must be too. Then for the second step, if we artificially stop frogs to enforce the rule in the non-backtracking model, then we can just get the self-similar model. Generally speaking, this further reduces root visits, which means if we can show that this restricted model has infinite visits, then the recurrence of the original frog model must hold.

Proposition 3. The non-backtracking and self-similar model can be coupled with the original frog model, such that every root visit in the simplified model corresponds to a distinct root visit in the original model.

Proof. After coupling the original frog model with non-backtracking rule, the new simplified path is a subset of the original path, and every root visit in the simplified path must appear in the original path at the same step, and thus **if the simplified model has infinitely many root visits, then so does the original model.** \square

Thus, proving that the simplified model visits the root an infinite amount of times is sufficient for proving the frog model is recurrent.

2.3 Generating function recursion

Definition. The *probability generating function* in this model is a mathematical tool that encodes the probability distribution of the number of visits V to the root in the frog model. Define the function $f : [0, 1] \rightarrow [0, 1]$ by

$$f(x) = \mathbb{E} [x^V] = \sum_{k=0}^{\infty} \mathbb{P}(V = k) x^k.$$

Here V denotes the number of visits to the root. And under this definition: $f(1) = \mathbb{P}(V < \infty)$. If the model is **recurrent** ($V = \infty$), then $f(1) = 0$.

Proposition 4. If $V = \infty$, then $f(x) = 0$ for $x \in [0, 1)$.

Proof. If $V = \infty$ with probability 1, then the original function becomes

$$f(x) = \mathbb{E}[x^\infty] = \sum_{k=0}^{\infty} \mathbb{P}(V = \infty) \cdot x^\infty,$$

for $x \in [0, 1)$, $x^\infty = 0$, so $f(x) = 0$. \square

Then the special case is when $x = 1$, as it maintains the continuity. We set $f(1) = \sum_{k=0}^{\infty} \mathbb{P}(V = k) = 0$, which implies if $V = \infty$, then $f \equiv 0$ in both cases.

Definition. A *functional operator* (denoted by A in this paper) is a mathematical object that takes a function as input and outputs another function by some computation. It acts as a ‘bridge’ that transforms the generating function.

We will now use a functional operator A to show that the probability generating function f satisfies the recursive distributional equation $f = Af$.

Proposition 5. Define A which is an operator on functions on $[0, 1]$ by

$$Ag(x) = \frac{x+2}{3} g\left(\frac{x+1}{2}\right)^2 + \frac{x+1}{3} g\left(\frac{x}{2}\right) \left(1 - g\left(\frac{x+1}{2}\right)\right).$$

The generating function f satisfies $f = Af$.

Proof. By the proposition above we conclude that if $V = \infty$, then $f(x) = 0$, and thus assume from now on that $\mathbb{P}[V = \infty] = 0 < 1$. Recall in the previous results we have that the subtrees have identical behavior as the original one. Then for the recursive structure of V , root visits can be decomposed into the visits from the initial active frog and the visits from awakened frogs in the subtrees. Each subtree’s contribution is an independent copy of V due to self-similarity. \square

To analyze the recursive structure of V , we define \emptyset as the origin and \emptyset' as its child. Let u and v be the children of \emptyset' with v will be visited first and then u . Due to the self-similar nature of the tree, the subtrees rooted at u and v are independent and identical copies of the full tree \mathbb{T}_2 . Define V_u and V_v to be the number of frogs subtrees $\mathbb{T}_2(u)$ and $\mathbb{T}_2(v)$ visit \emptyset' respectively. By self-similarity, both V_u and V_v are distributed identically to the original V , and also define V' which is distributed as V , and we can rewrite V as a composition of a pair of independent events:

$$V = 1\{\text{frog at } \emptyset' \text{ visit } \emptyset\} + 1\{u \text{ is visited}\} \text{Bin}\left(V', \frac{1}{2}\right) + \text{Bin}\left(V_v, \frac{1}{2}\right). \tag{2}$$

Now looking at the first three terms of equation (2), we can conclude that the first term accounts for

a possible visit to \emptyset by frog started at \emptyset' , the second term represents if u is visited, then the number of times that frog from $\mathbb{T}_2(u)$ visits \emptyset with binomial distribution $\text{Bin}(V', \frac{1}{2})$, and the third term is just the number of times the frog from $\mathbb{T}_2(v)$ that visit \emptyset also with Binomial distribution $\text{Bin}(V_v, \frac{1}{2})$.




Figure 3: Structure of V

Then we will firstly consider three disjoint events as A, B, C:

A: Initial frog moves to a child and activates its subtree (with probability $\frac{1}{3}$).

B: Initial frog doesn’t activate but another frog does (with probability $\frac{2(1-q)}{3}$).

C: There is no frogs that activate the other subtree (with probability $\frac{2q}{3}$, where $q = \mathbb{P}(\text{subtree isn’t activated})$).

When the initial frog activates both subtrees (Event A), we get terms involving products of generating functions, representing subtree contributions. And when there is only one subtree that gets activated (Event B), the expression involves conditional terms that account for this restricted activation. The precise form of A carefully weights these scenarios by their probabilities.




Figure 4: Events A, B and C

We can combine event A, B, C with equation (2), or in other words, recompose event A, B, C with conditions in three terms in equation (2).

Table 1: Contributions of Disjoint Events to Expected Root Visits

Event	Description	Probability	Conditional Expectation $\mathbb{E}[x^V \text{Event}]$
A	Activates both subtrees	1/3	$f\left(\frac{x+1}{2}\right)^2$
B	One subtree activated (restricted)	$2(1-q)/3$	$\frac{x+1}{2} \cdot \frac{f\left(\frac{x+1}{2}\right) - f\left(\frac{x}{2}\right)}{1-q} \cdot f\left(\frac{x+1}{2}\right)$
C	No activation of other subtree	$2q/3$	$\frac{x+1}{2} \cdot \frac{f(x/2)}{q}$

Let X' and X be the $\text{Bin}(V, \frac{1}{2})$ for term 2 and 3 respectively. Let Y be conditional Binomial distribution given V . Let Z be distributed as $\text{Bin}(V, \frac{1}{2})$ conditional on $\text{Bin}(V, \frac{1}{2}) = V$. Let $I \sim \text{Bernoulli}(\frac{1}{2})$, and since the three events A, B, C are independent of each other, we can then write V as $A \cup B \cup C$, or in other words,

$$V = \begin{cases} X' + X & \text{with probability } \frac{1}{3} \text{ by event A} \\ I + X' + Y & \text{with probability } \frac{2(1-q)}{3} \text{ by event B} \\ I + Z & \text{with probability } \frac{2q}{3} \text{ by event C} \end{cases}$$

where $q = 1 - p$ represents the probability that no frog ever visits the subtree rooted at node u .

Plug this into the expression of expectation of x^V and then we get:

$$\begin{aligned} \mathbb{E}(x^V) &= \frac{1}{3} \mathbb{E}(x^{X'+X}) + \frac{2(1-q)}{3} \mathbb{E}(x^{I+X'+Y}) \\ &\quad + \frac{2q}{3} \mathbb{E}(x^{I+Z}) \\ &= \frac{1}{3} \mathbb{E}(x^{X'}) \mathbb{E}(x^X) + \frac{2(1-q)}{3} \mathbb{E}(x^I) \mathbb{E}(x^{X'}) \\ &\quad \cdot \mathbb{E}(x^Y) + \frac{2q}{3} \mathbb{E}(x^I) \mathbb{E}(x^Z) \end{aligned}$$

To derive the operator A , we decompose the root visits V based on the movement of the initial frog and the subsequent of subtrees.

Let $q = \mathbb{P}(\text{subtree is not activated and } I \sim \text{Bernoulli}(1/2))$ represent the initial frog's potential return to the root. We partition the sample space into three disjoint events:

By the Law of Total Expectation,


$$f(x) = \sum_{i \in \{A, B, C\}} \mathbb{P}(i) \mathbb{E}[x^V | i].$$

Substituting the values from Table 1 below, the terms $1 - q$ and q in the denominators are canceled by the event probabilities, simplifying to:

$$\begin{aligned} f(x) &= \frac{1}{3} f\left(\frac{x+1}{2}\right)^2 + \frac{2(1-q)}{3} \cdot \frac{x+1}{2} \cdot \frac{f\left(\frac{x+1}{2}\right) - f\left(\frac{x}{2}\right)}{1-q} \\ &\quad \cdot f\left(\frac{x+1}{2}\right) + \frac{2q}{3} \cdot \frac{x+1}{2} \cdot \frac{f\left(\frac{x}{2}\right)}{q} \\ &= \frac{x+2}{3} f\left(\frac{x+1}{2}\right)^2 + \frac{x+1}{3} f\left(\frac{x}{2}\right) \left(1 - f\left(\frac{x+1}{2}\right)\right) \end{aligned}$$

Hence, we finally derive the operator A :

$$\begin{aligned} f(x) &= \frac{x+2}{3} f\left(\frac{x+1}{2}\right)^2 + \frac{x+1}{3} f\left(\frac{x}{2}\right) \\ &\quad \cdot \left(1 - f\left(\frac{x+1}{2}\right)\right) = Af(x) \end{aligned}$$



so $A^n f = f$.

Figure 5: $f = Af$

2.4 Proving recurrence

In this section, we will complete the proof of recurrence for the frog model on the binary tree ($d = 2$) by analyzing the generating function $f(x) = \mathbb{E}[x^V]$, where V is the number of visits to the root.

We will show that A is monotonic for functions on the set $S = \{g : [0, 1] \rightarrow [0, 1]\}$. We will then stochastically dominate f by a family of Poisson random variables, and apply the operator A . Finally we will show that $\lim_{n \rightarrow \infty} A^n 1 = 0$, meaning that $f = 0$. $f = 0$ implies the probability of finite visits to the root is 0. Therefore $V = \infty$ and the frog model is a.s recurrent.

Lemma 6. *The operator A is monotonic for functions belongs to the set $S = \{g : [0, 1] \rightarrow [0, 1], \text{nondecreasing}\}$.*

Proof. Let $g, h \in S$ and assume $g(x) \leq h(x)$ for all $x \in [0, 1]$. We analyze $Ag(x)$ term by term:

- First Term:** Since $g \leq h$ and both are non-decreasing, it follows immediately that $g\left(\frac{x+1}{2}\right)^2 \leq h\left(\frac{x+1}{2}\right)^2$.

2. **Second Term:** Consider the expression $\phi(g) = g(\frac{x}{2})[1 - g(\frac{x+1}{2})]$. Given $g \leq h$ and the recursive structure where $\frac{x}{2} \leq \frac{x+1}{2}$, the increase in the visit probability $g(x/2)$ dominates the potential decrease in the term $[1 - g]$, ensuring $\phi(g) \leq \phi(h)$.

Since both coefficients $\frac{x+2}{3}$ and $\frac{x+1}{3}$ are positive, the sum preserves the inequality. Thus, $Ag \leq Ah$, proving that A is monotonic on \mathcal{S} . \square

Lemma 7. *The operator A is invariant on \mathcal{S} , meaning if $g \in \mathcal{S}$, then $Ag \in \mathcal{S}$.*

Proof. To show $Ag \in \mathcal{S}$, we verify that $Ag(x)$ is both non-decreasing and bounded for $x \in [0, 1]$:

- **Non-decreasing:** $Ag(x)$ is a combination of terms $g(\frac{x+1}{2})^2$ and $g(\frac{x}{2})[1 - g(\frac{x+1}{2})]$. Since g is non-decreasing and the arguments $\frac{x}{2}, \frac{x+1}{2}$ are increasing functions of x , the sum (with positive coefficients) preserves the non-decreasing property.
- **Boundedness:** Given $g(x) \in [0, 1]$, each constituent term in the expression for $Ag(x)$ —including g^2 and the product $g(1 - g)$ —is bounded within $[0, 1]$. Furthermore, evaluating at the boundary $x = 1$ yields $Ag(1) = 1$ (assuming $g(1) = 1$), confirming $Ag(x) \in [0, 1]$ for all x .

Thus, A maps \mathcal{S} into itself. \square

We will now use a *Poisson* of any mean to stochastically dominate the frog model. We will examine how A acts on a family of *Poisson* random variables a .

Remark. Define $g_a(x) = e^{a(x-1)}$ as the *Poisson* generating function for variables a . We will omit the proof of the following result as outside the scope of the paper: For all $a \geq 0$, then for all $x \in [0, 1]$, there will be $g_a(x) \leq g_{a+c_a}(x)$, where

$$c_a = \begin{cases} \frac{1}{3}e^{-2} & 0 \leq a \leq 4, \\ \frac{1}{3}e^{-\frac{a}{2}} & a \geq 4. \end{cases}$$

Lemma 8. *For all $x \in [0, 1]$, $\lim_{n \rightarrow \infty} A^n g_0(x) = 0$.*

Proof. Let $g_a(x) = e^{a(x-1)}$ and $a_{n+1} = a_n + c_{a_n}$ with $a_0 = 0$. Note $a_n \rightarrow \infty$ since $c_a > 0$. We show $A^n g_0 \leq g_{a_n}$ by induction. The base case ($n = 0$) holds as $A^0 g_0 = 1 = g_{a_0}$. Assuming the hypothesis for $n = k$, the monotonicity of A and the property $g_a \leq g_{a+c_a}$ imply:

$$A^{k+1} g_0 \leq A g_{a_k} \leq g_{a_k+c_{a_k}} = g_{a_{k+1}} \tag{3}$$

Thus, $0 \leq \lim_{n \rightarrow \infty} A^n g_0(x) \leq \lim_{n \rightarrow \infty} e^{a_n(x-1)} = 0$ for $x < 1$. By the Squeeze Theorem, the limit is 0. \square

By definition since f is a probability generating function, $f \leq 1 = g_0(x)$ for $x \in [0, 1]$. We know that A is monotonic, and have that $A^n f = f \leq A^n 1$. Since $\lim_{n \rightarrow \infty} A^n 1 = 0$, f must be constantly 0. Therefore V is a.s. infinite, and by the coupling of the number of visits to the root of the frog model is also infinite.

Therefore we have proved the recurrence for the frog model on the binary tree, which give a strong evidence for **Theorem 1**: The frog model is recurrent when $d \leq 2$.

3 CONCLUSION

In this paper, we explored the recurrence of the frog model on d -ary trees by using the concept of stochastic dominance. The frog model was a very interesting model of self-interacting random walks to explore. It has applications on the spread of information and disease that can hopefully be further investigated in the future. The usage of a simpler model to stochastically dominate the frog models was vital in proving its recurrence. In the whole process, we use recursive distributional equations and generating functions to simplify stochastic systems.

REFERENCES

- [1] Ravi P. Agarwal, Maria Meehan, Donal O'Regan "Fixed Point Theory and Applications" Cambridge University Press, Chapter 4, pp. 85–89, (2001)
- [2] Nicholson, W. K. "Linear Algebra with Applications", Chapter 9.3 (2019)
- [3] Christopher Hoffman, Tobias Johnson, Matthew Junge "Recurrence and transience for the frog model on trees," The Annals of Probability, Ann. Probab. 45(5), 2826-2854, (September 2017)
- [4] Alves, O. S. M., Machado, F. P. and Popov, S. Y. "Phase transition for the frog model." Electron. Journal of Probability 7(16), 1–21, (2002)
- [5] Pemantle, R. "A survey of random processes with reinforcement." Probability Surveys. 4(1), 1–79, (2007)
- [6] Telcs, A. and Wormald, N. C. "Branching and tree indexed random walks on fractals." Journal of Applied Probability 36(4), 999–1011, (1999)

INTERVIEW WITH ANUSH TSERUNYAN

Louis Meunier and Nitya Khirwar



$\delta\epsilon$: Could you briefly introduce yourself; your academic position, and where you are from?

Hi, my name is Anush Tserunyan. I'm an associate professor at the Department of Mathematics and Statistics at McGill, and I'm originally from Armenia. I did my PhD and postdoc at UCLA, and got my first associate professorship at the University of Illinois Urbana-Champaign. Then, I moved to McGill in 2020 during COVID, and became an associate professor here in 2023.

$\delta\epsilon$: Could you describe your research interests in a general sense, as well as a specific question you're interested in now?

My research is rooted in a branch of logic called descriptive set theory (DST), which really has very little to do with logic or set theory, but is rather a blend of analysis and combinatorics. The questions that I usually work on are a blend of dynamics — namely measurable dynamics and Ergodic theory — and measured group theory. Both explore the actions of groups on probability spaces, but while dynamics explores the actions, group theory helps use them to study the group. Since the 80s and 90s, DST has focused on studying equivalence relations and graphs on Polish spaces, such as the space of Borel sets. DST provides means

for ensuring that when you deal with these equivalence relations, you define more sets and functions so that they remain Borel. DST has these uniformization theorems that help us ensure that equivalence relations on, say, Borel sets remain Borel. A theorem that I use every day before brushing my teeth is called the Lusin-Novikov theorem, which says that if you have a Borel equivalence relation whose classes are countable, like a countable group that acts on the reals and the action maps Borel sets to Borel sets, then the equivalence relation you get from this action is a Borel equivalence relation with countable classes that are each as big as the group. Turns out, all such countable Borel equivalence relations arise from this.

There are equivalence relations, like the ones with countable classes and Borel sets, which are called “treeable”: these are the ones that allow acyclic graphs whose connected components are exactly those classes of equivalence relations. The caveat here is that the graph has to be Borel, otherwise its existence is trivial by the axiom of choice: you choose a point from every graph and draw a star, but this graph would be useless, because it's not Borel or measurable with respect to any measure. I conjecture, and some other people did too independently, that if an equivalence relation admits such a graphing and if this graph excludes some finite minor then you can replace them with a Borel acyclic graph, which are trees. So in other words, the equivalence relation that admits a minor excluding graph must be treeable. So there is some work going on right now with me, my current and former PhD students and visiting scholars.

$\delta\epsilon$: On the day-to-day, what does your work with collaborators look like on this topic?

It's meetings after meetings, if that's what you're asking. This semester I'm teaching two classes so I'm drowning, but normally, like last semester, even on the days of teaching, I had research meetings. This semester, I had to reduce them to be able to stay alive. So on, say, Monday, I have three research meetings. And some of them can go on for hours, and they're usually on Zoom because most of my collaborators aren't around. I collaborate with the founder of a big part of my research area in France, Damien Gaboriau, and with my academic twin brother, which means we had the same advisor and graduated in the same year, so we're twins. So I collaborate with him a lot, he's at the University of Florida. And then I collaborate with these PhD students that are graduating and are visiting

me. It's really cool to have to talk to other people who are as enthusiastic about the stuff as I am, and have more fresh ideas. Like, probabilistic methods have become very important for my area, and I have a paper using probability like percolation, but I'm still a beginner in this field.

$\delta\epsilon$: Switching topics a bit, you've studied in Yerevan State University and UCLA, and taught at UIUC and McGill. What are the biggest cultural differences you've observed in how math is approached or taught across these universities?

In my undergrad in Armenia, the things we were taught were ancient, and it was all following Soviet textbooks. Everything was so complicated that afterwards when I learned, say, what the metric space is, I realized that 80% of this book could be done in 3 pages in baby Rudin. There wasn't much motivation for the things we were learning. But right now I'm actually teaching there. I do one honours course per year voluntarily, for kids hopefully like me or better than me that are there and they otherwise would just become programmers or something, but they're good at math and they maybe would want to continue in math but they don't see prospects in it and go into industry because of that. And even if they go into industry, it would be good if we make some killer whales and send them to industry, it would be also good for my country.

But even in these courses that I teach — I've been teaching three years now — I have had no woman in Armenia take my course. [...] But we have it amazing here in Canada. In Armenia, they'd say if a woman gets a 100 out of a 100 that she's just working hard, but that score doesn't come from hard work. It's ingenuity. Here, the top of my class are often women. I see stronger women day by day, and McGill has the strongest undergraduate population I've seen in my career. This culture of summer research, the culture where undergrads need to take graduate-level courses, I think that's a really good thing to have. I also have some undergrads who come to the research seminars, I hadn't seen anything like this before McGill.

$\delta\epsilon$: At McGill specifically, is there a particular area of math or subject that isn't emphasized enough or maybe is?

I mean, I was talking to a student recently, who said he wants to learn more analytic number theory, but there is no course offered like this. Algebraic number theory, we have courses, but analytic number theory, we don't because there are no analytic number theorists in this department. There's ergodic theory, but it's only basically me who does it, maybe Marcin [Sabok] in a little bit, but like we are fake ergodic theorists. We're logicians that turned into ergodic theorists. It's not just

a McGill problem though, we don't have any in Quebec. Also, the only logic course here is [MATH] 318. It's not an honours course, and it's just an elective or something. And it's maybe mandatory for CS students, so most of it is CS students. I feel like it's a really strong class and very difficult. So this should be turned into a course that honours students would also take and would get credit for. Yeah, it's only, you don't know what isn't there, because you only know what is there. In that sense, as I said, I think you guys have more courses and [they're] deeper than the schools I know.

$\delta\epsilon$: You were recently mentioned in a Quanta magazine about new work in descriptive set theory. If you read the article, how do you think it did at conveying its subject to a general audience?

I like to read their articles, but not the ones in my area, because those ones I can detect when they're badly written. I could not finish reading that article. Some things were more dramatic than they should have been, and then some things were not dramatic enough. Also, things were wishy-washy; there were false statements. In a different subject, I wouldn't mind such a level of wishy-washiness, but in my subject, it's like — can you watch a movie about mathematicians?

$\delta\epsilon$: Do you have any advice for current math undergraduates?

So to those who see themselves as continuing in grad school, I would say learn as much as you can and as deep as you can. So don't take a million courses and learn everything on the surface. Most of my firm knowledge comes from what I learned in undergrad, and because my undergrad was in CS, really I only learned math in my PhD. So whatever I learned in courses in my PhD is the knowledge I have. Everything else that I learned afterwards is surface-level. So this is your chance to learn.

[...] Fight every definition, fight every theorem, so that when I ask you, "Why was this hypothesis put there, couldn't we get away without it?" or "why are we doing this subject anyway? Why are we studying linear algebra?" or "why are we studying analysis? Why is a Cauchy sequence important?" You know, try to answer these kinds of questions and ask them to yourself. Basically, try to learn as much as deeply as you can and as much as you can. But when you learn, try to also fight it as much as you can, so that if I come and try to bully you in the subjects, you actually withstand me. And ask questions like, why is this subject important? Why should we care? Students in topology ask me every day, "Why do we care about point set topology? Why aren't metric spaces good enough?" [...] Try to understand why. So yeah, this is my advice.

OPTIMAL SCORING STRATEGIES FOR THE LIZARD CULT

Cedric Phillips and Adelaide Row

We propose a discrete-time Markov chain framework to model scoring using gardens for the Lizard Cult faction in the asymmetric board game *Root*. The Cult’s ability to score from their gardens is affected by three different processes with stochasticity introduced by their opponents: suited card draws from a shared deck, the dynamic Outcast suit determined by recently discarded cards, and forcible removal of their gardens. Unlike other factions, the Cult chooses when it will score from its gardens, which allows for significant strategy-based variation. We use Monte Carlo simulation to visualize scoring trajectories and determine which parameters are most key.

1 INTRODUCTION

Root is a board game with a high degree of asymmetry: each player chooses a unique faction out of 10. Every faction has unique actions and their own ways to score and ‘win the war’. For the Lizard Cult faction, the primary victory mechanism involves building *gardens* and performing the Score ritual “citeroot” (Appendix F). Most other factions have automatic or incidental automatic scoring, but the Cult chooses when it scores points (at the expense of its resources), which creates an optimization problem. Three sources of randomness complicate this optimization:

1. **Card draws:** Rituals require suited cards drawn from a shared 54-card deck. The Cult’s ‘Hatred of Birds’ rule means the 18 Bird-suit cards cannot be used for scoring.
2. **Garden dynamics:** Gardens grow in number through player actions (the Build ritual, the Cult’s Conspiracies) but are likely to be destroyed by opponents through combat and special abilities.
3. **Hand composition:** The random distribution of which suits of cards will be drawn cards constrains the suits which can be scored each turn.

1.1 Game Mechanics

The Cult’s faction board has three garden tracks (Fox, Rabbit, Mouse), each with capacity 5. Let $\mathbf{G} = (G_F, G_R, G_M) \in \{0, \dots, 5\}^3$ denote gardens on the map. The Cult starts with 1 garden (of chosen suit) on the map “citeroot”.

Scoring. On the Cult’s turn, performing the Score ritual for suit s yields victory points based on garden

count, as per the faction board:

$$S(g) = \begin{cases} 0 & g \in \{0, 1\} \\ 2 & g \in \{2, 3\} \\ 3 & g = 4 \\ 4 & g = 5 \end{cases} \quad (1)$$

Card draw. The Cult draws $n(\mathbf{G}) = 1 + \text{Bonus}(\mathbf{G})$ cards per turn, where

$$\text{Bonus}(\mathbf{G}) = \sum_s \mathbb{1}[G_s \geq 2], \quad (2)$$

counts draw bonuses (one per suit with ≥ 2 gardens). After drawing, the Cult must discard down to a hand limit of 5 cards. Since discards are chosen by the player, this allows preferential retention of suited cards for scoring on subsequent turns.

Scoring constraints. Each suit can be scored at most once per turn, but the Cult can score multiple suits (up to 3) in a single turn. Scoring requires spending a suited card.

Victory. The first player to 30 VP wins immediately.

1.2 Threshold Policies

We formalize the Cult’s choice as to whether they should score or wait as a threshold policy.

Definition (Threshold Policy). *Our threshold policy π_θ with parameter $\theta \in \{1, \dots, 5\}$ is:*

- **Score** suit s if $G_s \geq \theta$ and a suit- s card is available
- **Build**¹ otherwise

Definition (Victory Time). *The victory time under policy π_θ is the hitting time*

$$\tau(\theta) = \inf\{t \geq 0 : V_t \geq 30\} \quad (3)$$

and we define $\theta^* = \arg \min_\theta \mathbb{E}[\tau(\theta)]$, the minimizer of expected turns to victory.

¹The mechanisms that ‘Build’ encompasses are vast, and we cannot model them all here. In general, the Cult should use the Sacrifice ritual on Bird cards, and complete Recruit rituals. These will complement the Build ritual and Sanctify conspiracy, which allow the Cult to get more gardens on the map.

Standard Root Map (Fall)

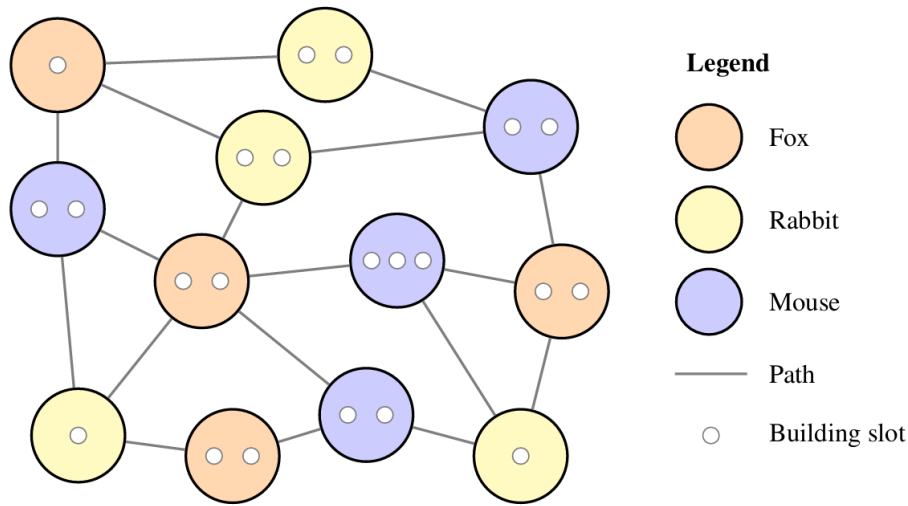


Figure 1: The standard (Fall) map. Clearings are coloured by suit; white circles indicate all possible building slots where the Lizard Cult can place gardens (note: in some games, certain slots have Ruins and are unusable) [1].

2 CARD DRAW ANALYSIS

2.1 The Hypergeometric Distribution

The deck contains 54 cards: 12 Fox, 12 Rabbit, 12 Mouse, 18 Bird. The Cult’s ‘Hatred of Birds’ means only 36 cards (two-thirds) are usable for scoring.

Proposition 1 (Draw Distribution). *Drawing n cards, the number of suit- s cards follows*

$$X_s \sim \text{Hypergeometric}(54, 12, n) \tag{4}$$

with mean $\mathbb{E}[X_s] = 12n/54 = 2n/9$.

Proof. This is a standard hypergeometric distribution with parameters making it mathematically equivalent to drawing cards from a deck without replacement [4]. The mean follows from linearity of expectation. \square

Proposition 2 (Hit Probability). *The probability of drawing at least one card of suit s is*

$$q(n) = \mathbb{P}(X_s \geq 1) = 1 - \frac{\binom{42}{n}}{\binom{54}{n}}. \tag{5}$$

Since the Cult can score multiple suits per turn, the expected VP per turn depends not just on the per-suit hit probability $q(n)$ but on how many distinct suits appear in a single draw. We therefore consider $K = |\{s : X_s \geq 1\}|$, the number of scorable suits (i.e., those with at least one available card).

Cards drawn n	1	2	3	4
$q(n)$	0.222	0.399	0.540	0.652
$\mathbb{E}[\text{suited cards}]$	0.67	1.33	2.00	2.67
$\mathbb{E}[K]$	0.67	1.20	1.62	2.20

Furthermore, for $n = 4$ draws from a fresh deck:

$K \geq$	1	2	3
Probability	0.990	0.759	0.418

In particular, with full draw bonuses ($n = 4$), the Cult has 76% chance of scoring 2+ suits per turn.

Proof. See (Appendix B). \square

3 MARKOV CHAIN MODEL

We define a finite-state Markov chain that captures the essence of the relevant strategic working of the Cult.

3.1 Reduced State Space

Tracking the full game state is both overly complicated and quite difficult; the state space becomes intractably large. We instead define a reduced state space for the key aspects of strategy.

Definition (Reduced State Space). *The state on turn t for the state space \mathcal{S} is given by*

$$\mathbf{X}_t = (\mathbf{G}_t, V_t, \mathbf{h}_t) \in \mathcal{S} \tag{6}$$

where:

- $\mathbf{G}_t = (G_F^t, G_R^t, G_M^t) \in \{0, 1, 2, 3, 4, 5\}^3$: garden counts per suit
- $V_t \in \{0, 1, \dots, 30\}$: victory points (states with $V \geq 30$ are absorbing)
- $\mathbf{h}_t = (h_F^t, h_R^t, h_M^t) \in \{0, 1\}^3$: indicators where $h_s = 1$ iff at least one suit- s card is in hand.

The garden count is given by $|\mathbf{G}| = G_F + G_R + G_M$. The state space has $|\mathcal{S}| = 6^3 \times 30 \times 2^3 = 51,840$ transient states, and one absorbing ‘victory’ state.

The hand indicator \mathbf{h} captures whether scoring suit s is possible this turn, instead of considering how many suited cards are available. This does not affect the analysis, as the Cult can only score each suit once per turn. The other approximation is a fresh deck: in the model, we draw from the full 54-card deck each turn. In particular, draws from the deck are independent of existing hand composition. This simplification is necessary for computation (otherwise, there are 41,743 deck states) and reasonable since the Cult only cares about suit distribution (and not specific cards).

3.2 Turn Structure and Transition Kernel

We adapt *Root*’s turn structure language of *Birdsong*, *Daylight*, and *Evening*. In *Birdsong*, the Cult can Sanctify new gardens into play; in *Daylight*, they can perform rituals, including Build and Score. Logically, we assume that players will Build before they Score, allowing us to simplify. They then Draw, and pass to other players. We model this in four phases: growth, scoring, card draw, then attrition.

$$\begin{aligned} (\mathbf{G}, V, \mathbf{h}) &\xrightarrow{1: \text{Growth}} \mathbf{G}' \xrightarrow{2: \text{Score}} (V', \mathbf{h}') \\ &\xrightarrow{3: \text{Draw}} \mathbf{h}'' \xrightarrow{4: \text{Attrition}} (\mathbf{G}'', \mathbf{h}'''). \end{aligned}$$

The next-turn state is $\mathbf{X}_{t+1} = (\mathbf{G}'', V', \mathbf{h}''')$.

3.2.1 Phase 1: Growth of Gardens

Gardens can grow through Build rituals and the Sanctify conspiracy; since this depends highly on player behaviour, luck, and other factors, we choose to model this as independent Bernoulli growth using an aggregate parameter μ .

Definition (Growth Transition). *Each suit’s garden count increases by 1 with probability μ , capped at 5:*

$$\begin{aligned} \mathbb{P}(G'_s = g+1 \mid G_s = g) &= \mu \cdot \mathbb{1}[g < 5] \\ \mathbb{P}(G'_s = g \mid G_s = g) &= 1 - \mu \cdot \mathbb{1}[g < 5] \end{aligned} \quad (7)$$

Suits grow independently, so

$$\mathbb{P}(\mathbf{G}' \mid \mathbf{G}) = \prod_{s \in \{F, R, M\}} \mathbb{P}(G'_s \mid G_s). \quad (8)$$

Remark (Interpretation of μ). In practice, $\mu \in [0.2, 0.5]$ depending on luck and opponent pressure (for example, the Cult may need to defend gardens instead of building new ones).

3.2.2 Phase 2: Scoring

Under a threshold policy π_θ , the Cult scores all eligible suits.

Definition (Scoring Transition). *Given threshold θ , garden configuration \mathbf{G}' , and hand \mathbf{h} :*

- For each suit s : if $G'_s \geq \theta$ and $h_s = 1$, score $S(G'_s)$ points and set $h'_s = 0$.
- Otherwise: $h'_s = h_s$.

The VP update is

$$V' = \min \left(V + \sum_s S(G'_s) \cdot \mathbb{1}[G'_s \geq \theta] \cdot h'_s, 30 \right). \quad (9)$$

The hand update is

$$h'_s = h_s \cdot (1 - \mathbb{1}[G'_s \geq \theta]). \quad (10)$$

3.2.3 Phase 3: Card Draw and Hand Update

The Cult draws cards based on draw bonuses, then updates hand composition.

Definition (Draw Count). *Given garden configuration \mathbf{G}' , the number of cards drawn is*

$$n(\mathbf{G}') = 1 + \text{Bonus}(\mathbf{G}') = 1 + \sum_s \mathbb{1}[G'_s \geq 2]. \quad (11)$$

Definition (Hand Transition). *Let $\mathbf{Y} = (Y_F, Y_R, Y_M, Y_B)$ be the cards drawn, where Y_S is the count of suit- s cards, and denote $n = n(\mathbf{G}')$. Under a fresh-deck approximation, \mathbf{Y} follows a multivariate hypergeometric distribution:*

$$\mathbf{Y} \sim \text{MultiHypergeom}(54; 12, 12, 12, 18; n). \quad (12)$$

After optimal discard, the new hand indicator is

$$h''_s = \mathbb{1}[h_s = 1 \text{ or } Y_s \geq 1]. \quad (13)$$

Since we only need to know whether $Y_s \geq 1$, the transition probability simplifies to

$$\mathbb{P}(h''_s = 1 \mid h_s, n) = \begin{cases} 1 & \text{if } h_s = 1 \\ q(n) & \text{if } h_s = 0 \end{cases} \quad (14)$$

where $q(n)$ is the probability of drawing at least one card of suit s (see Proposition 2).

Due to the fresh-deck approximation, drawing the suits are conditionally independent:

$$\mathbb{P}(\mathbf{h}'' \mid \mathbf{h}', \mathbf{G}') = \prod_s \mathbb{P}(h''_s \mid h_s, n(\mathbf{G}')). \quad (15)$$

Definition (Transition Matrix). We construct the transition matrix \mathbf{P}_θ , a $(|\mathcal{S}| + 1) \times (|\mathcal{S}| + 1)$ matrix with entries $[\mathbf{P}_\theta]_{\mathbf{x}, \mathbf{x}'} = \mathbb{P}_\theta(\mathbf{x}' | \mathbf{x})$. States with $V' \geq 30$ transition to the absorbing victory state with probability 1. Let

$$\mathbf{P}_\theta = \begin{pmatrix} \mathbf{Q}_\theta & \mathbf{r}_\theta \\ \mathbf{0}^\top & 1 \end{pmatrix} \quad (20)$$

where \mathbf{Q}_θ is the $|\mathcal{S}| \times |\mathcal{S}|$ submatrix of transient-to-transient transitions, $\mathbf{0}$ is the zero vector of dimension $|\mathcal{S}|$, and $\mathbf{r}_\theta \in [0, 1]^{|\mathcal{S}|}$ is the absorption vector whose \mathbf{x} -th entry gives the probability of transitioning from transient state \mathbf{x} directly to the absorbing victory state in one step:

$$[\mathbf{r}_\theta]_{\mathbf{x}} = \mathbb{P}_\theta(V' \geq 30 | \mathbf{X}_t = \mathbf{x}). \quad (21)$$

Note that absorption can only occur during the scoring phase (Phase 2); the subsequent draw and attrition phases cannot increase VP.

3.4 Victory Time Analysis

Definition (Victory Time). The victory time under policy π_θ starting from state \mathbf{x} is

$$\tau_\theta(\mathbf{x}) = \inf\{t \geq 0 : V_t \geq 30 | \mathbf{X}_0 = \mathbf{x}\}. \quad (22)$$

Proposition 4 (Hitting Time Moments). Let $\mathbf{N}_\theta = (\mathbf{I} - \mathbf{Q}_\theta)^{-1}$ be the fundamental matrix, and $\mathbf{1}$ be the all-ones vector of dimension $|\mathcal{S}|$. Then:

1. The expected victory time from state \mathbf{x} is

$$\mathbb{E}[\tau_\theta(\mathbf{x})] = [\mathbf{N}_\theta \cdot \mathbf{1}]_{\mathbf{x}}. \quad (23)$$

2. The variance is

$$\text{Var}(\tau_\theta(\mathbf{x})) = [(2\mathbf{N}_\theta - \mathbf{I}) \cdot \mathbf{t}_\theta]_{\mathbf{x}} - [\mathbf{t}_\theta]_{\mathbf{x}}^2. \quad (24)$$

for $\mathbf{t}_\theta = \mathbf{N}_\theta \cdot \mathbf{1}$ the expected hitting time vector.

3. The probability of victory in exactly k turns is

$$\mathbb{P}(\tau_\theta(\mathbf{x}) = k) = [\mathbf{Q}_\theta^{k-1} \cdot \mathbf{r}_\theta]_{\mathbf{x}}. \quad (25)$$

Proof. These are standard results for absorbing Markov chains. The matrix $\mathbf{N} = (\mathbf{I} - \mathbf{Q})^{-1}$'s role in computing expected absorption times is developed in Kemeny and Snell [3], Chapter 3, Section 3. \square

Definition (Initial State). The Cult begins with one garden in a chosen clearing. Without loss of generality, we fix this as Fox, giving $\mathbf{G}_0 = (1, 0, 0)$ and $V_0 = 0$. The Cult also draws 3 cards from the full 54-card deck, so the initial hand state \mathbf{h}_0 is random. Under the fresh-deck approximation, the initial distribution is

$$\pi_0(\mathbf{h}) = \prod_{s \in \{F, R, M\}} q(3)^{h_s} (1 - q(3))^{1 - h_s} \quad (26)$$

where $q(3) \approx 0.537$ is the probability of drawing at least one card of a given suit in 3 draws (c.f. Proposition 2). The expected victory time from this initial distribution is

$$\mathbb{E}[\tau_\theta] = \sum_{\mathbf{h} \in \{0, 1\}^3} \pi_0(\mathbf{h}) \cdot [\mathbf{t}_\theta]_{((1, 0, 0), 0, \mathbf{h})} \quad (27)$$

where $\mathbf{t}_\theta = (\mathbf{I} - \mathbf{Q}_\theta)^{-1} \mathbf{1}$ is the hitting time vector.

The transition matrix \mathbf{P}_θ is sparse and can be easily constructed; the algorithm is in (Appendix A).

4 OPTIMAL THRESHOLD DERIVATION

We prove that $\theta^* = 2$ is optimal via two approaches: delay value analysis and Markov chain computation.

4.1 Delay Value Characterization

We calculate the expected VP change from waiting an additional turn before scoring a suit.

Definition (Delay Value). For a suit with g gardens, the delay value is

$$\Delta(g) = \mathbb{E}[\text{VP delay-then-score}] - \mathbb{E}[\text{VP score-now}].$$

Immediate scoring is optimal when $\Delta(g) \leq 0$.

Let μ be the growth probability and $p_\ell = 1 - e^{-\lambda/|G|}$ be the per-suit loss probability ($= 1 - e^{-\lambda/3}$ under diversified play). If we delay one turn, the garden count evolves to $g' \in \{g-1, g, g+1\}$: with probability $\mu(1 - p_\ell)$ the garden grows and survives, with probability $(1 - \mu)(1 - p_\ell)$ nothing changes, and with probability p_ℓ a garden is lost (assuming that attrition removes at most one garden per suit, valid for small λ). Then

$$\begin{aligned} \mathbb{E}[V_{\text{delay}}] &= \mu(1 - p_\ell)S(g+1) + (1 - \mu)(1 - p_\ell)S(g) \\ &\quad + p_\ell S(g-1) \\ \implies \Delta(g) &= \mu(1 - p_\ell)[S(g+1) - S(g)] \\ &\quad - p_\ell[S(g) - S(g-1)]. \end{aligned}$$

Evaluating at each garden count with parameters $\mu = 0.4$, $p_\ell = 0.065$ (corresp. to $\mu = 0.4$, $\lambda = 0.2$):

g	1	2	3	4	5
$S(g)$	0	2	2	3	4
Gain	2	0	1	1	—
Loss	0	2	0	1	1
$\Delta(g)$	+0.75	-0.13	+0.37	+0.31	-0.065

Two features stand out. At $g = 1$, scoring yields 0 VP, so there is no opportunity cost to waiting; $\Delta(1) > 0$ for all parameter values. At $g = 2$, the scoring plateau $S(2) = S(3) = 2$ means growth provides zero upside, but attrition to $g = 1$ destroys all scoring value (downside = 2). In particular, $\Delta(2) = -2p_\ell < 0 \forall \mu$.

4.2 Optimal Threshold

We show that $\theta^* = 2$ by considering the build-up time. In particular, for $\mu \in [0.2, 0.6]$ and $\lambda \in [0.1, 0.5]$, the optimal threshold is $\theta^* = 2$.

We provide some reasoning for this idea, but not a formal proof.

Claim 5. ($\theta = 1$ is dominated). *Since $S(1) = 0$, a policy that scores at $g = 1$ earns zero VP from any suit with exactly one garden. Since $S(g) \geq 2$ for all $g \geq 2$, any threshold $\theta \geq 2$ dominates $\theta = 1$ always.*

Claim 6. (*Immediate scoring is optimal at $g = 2$). As shown in the table above, $\Delta(2) = -2p_\ell < 0$ for all $\lambda > 0$. This reflects the structure of S : the plateau $S(2) = S(3) = 2$ means growth from $g = 2$ to $g = 3$ provides zero VP, but attrition from $g = 2$ to $g = 1$ causes a total loss of scoring value.*

Claim 7. ($\theta \geq 3$ is suboptimal). *Although $\Delta(3)$ and $\Delta(4)$ are positive, a threshold policy must first build to that level before scoring begins.*

For $\theta = 3$, the case is immediate: since $S(3) = S(2)$ and the draw bonus activates at $g = 2$, the scoring rate is identical to $\theta = 2$'s, and $\theta = 3$ requires strictly more build-up time so $\mathbb{E}[\tau(3)] > \mathbb{E}[\tau(2)] \forall \mu, \lambda$.

For $\theta \geq 4$, we use a decomposition (Appendix C): $\mathbb{E}[\tau(\theta)] \approx T_{\text{build}}(\theta) + 30/\bar{V}(\theta)$, where $\bar{V}(\theta) = S(\theta) \cdot \mathbb{E}[K | n(\theta)] \cdot p_{\text{surv}}$ is the expected VP per scoring turn.

Since all thresholds $\theta \geq 2$ yield $n(\theta) = 4$ (three suits above the draw bonus threshold), $\mathbb{E}[K | n]$ and p_{surv} are common factors. The saving from raising the threshold to θ is therefore

$$\Delta_{\text{score}}(\theta) = \frac{30}{S(2) \cdot C} - \frac{30}{S(\theta) \cdot C} = \frac{30}{C} \cdot \frac{S(\theta) - S(2)}{S(2) \cdot S(\theta)}$$

where $C = \mathbb{E}[K | n=4] \cdot p_{\text{surv}}$ is independent of θ . Since $C \leq 3$ (at most 3 suits scored) and $S(\theta) \leq 4$,

$$\Delta_{\text{score}}(\theta) \leq \frac{30}{C} \cdot \frac{S(\theta) - 2}{2 \cdot S(\theta)} \leq \frac{30}{C} \cdot \frac{1}{4} = \frac{7.5}{C} \quad (28)$$

using $S(5) = 4$, the largest available value. Since $C \geq 1$, this gives $\Delta_{\text{score}} \leq 7.5$ turns saved.

Meanwhile, the additional build-up cost satisfies

$$\Delta_{\text{build}}(\theta) = T_{\text{build}}(\theta) - T_{\text{build}}(2) \geq \frac{\theta - 2}{\mu - \lambda/3} \quad (29)$$

since the net per-suit drift is at most $\mu - \lambda/3$, and all three suits must reach θ . For the parameter range $\mu \in [0.2, 0.6]$ and $\lambda \in [0.1, 0.5]$, the net drift $\mu - \lambda/3$ ranges from 0.03 to 0.57, giving:

θ	$\Delta_{\text{score}} \leq$	$\Delta_{\text{build}} \geq$	Net
4	$\frac{30}{C} \cdot \frac{1}{6} \approx 2.5$	3.5 to 67	loss
5	$\frac{30}{C} \cdot \frac{1}{4} \approx 3.8$	5.3 to 100	loss

(with the most favorable value $C = 2$). Under the most favorable reasonable drift ($\mu = 0.6$, $\lambda = 0.1$), Δ_{build} exceeds Δ_{score} for $\theta \geq 4$, so in the viable regime $\mu > \lambda/3$, the build penalty dominates.

4.3 Endgame Adjustment

When $V + S(G_s) \geq 30$, the optimal action is to and score immediately and win, regardless of θ .

Corollary 8 (Victory-Aware Policy). *The optimal threshold is state-dependent:*

$$\theta^*(V) = \min(\theta_{\text{base}}^*, g^* \text{ s.t. } S(g^*) \geq 30 - V). \quad (30)$$

4.4 Exact Markov Chain Computation

The transition matrix \mathbf{P}_θ defined in Section 3 enables exact computation of $\mathbb{E}[\tau(\theta)]$.

Remark (Exact Victory Times). Let π_0 be the initial distribution from Definition 3.4, with $\mathbf{G}_0 = (1, 0, 0)$ and \mathbf{h}_0 drawn from 3 initial cards. The fundamental matrix $\mathbf{N}_\theta = (\mathbf{I} - \mathbf{Q}_\theta)^{-1}$ yields the hitting time vector $\mathbf{t}_\theta = \mathbf{N}_\theta \cdot \mathbf{1}$, and the expected victory time is $\mathbb{E}[\tau_\theta] = \pi_0 \cdot \mathbf{t}_\theta$ (equation (27)). Solving the sparse linear system $(\mathbf{I} - \mathbf{Q}_\theta)\mathbf{t} = \mathbf{1}$ gives exact values for these quantities.

θ	1	2	3	4	5
$\mathbb{E}[\tau_\theta]$	∞	10.95	12.47	13.85	16.03
$\text{Var}(\tau_\theta)$	—	6.68	9.04	11.31	14.91

Parameters: $\mu = 0.4$, $\lambda = 0.2$. Expectations taken over both the initial hand draw and the trajectory.

5 DIVERSIFICATION ANALYSIS

The threshold policy π_θ prescribes *when* to score, but the Cult must also decide *how to build*: should gardens be spread evenly across suits or concentrated into one? The Markov chain model provides an answer.

Under $\theta^* = 2$, a garden at $g = 1$ contributes nothing, and a garden at $g = 3$ contributes no more than $g = 2$ since $S(3) = S(2)$. The strategy question is therefore whether to activate all three scoring tracks at $g = 2$ or to push a suit higher at the expense of another.

Proposition 9 (Diversification Dominance). *Let $\mathbf{G}_D = (2, 2, 2)$ and $\mathbf{G}_C = (4, 2, 0)$ be two configurations with 6 total gardens. Then*

$$\mathbb{E}[V_{\text{turn}} | \mathbf{G}_D] > \mathbb{E}[V_{\text{turn}} | \mathbf{G}_C]. \quad (31)$$

Proof. (Sketch). The configurations differ in two quantities from our model. First, \mathbf{G}_D activates all three draw bonuses ($n = 4$), while \mathbf{G}_C activates only two ($n = 3$). By Proposition 2, the per-suit hit probability increases from $q(3) = 0.540$ to $q(4) = 0.652$.

Second, \mathbf{G}_D opens three scoring tracks (each worth $S(2) = 2$), while \mathbf{G}_C opens only two (worth $S(4) + S(2) = 3 + 2 = 5$ combined).

Furthermore, at $\lambda = 0.2$, the per-suit survival probability for the diversified strategy is $p_{2,\text{surv}} = e^{-\lambda \frac{2}{6}} \approx 0.936$, the same as the concentrated strategy in its 2-garden suit, but in the 4-garden suit it is $p_{4,\text{surv}} = e^{-\lambda \frac{4}{6}} \approx 0.875$. Therefore:

$$\begin{aligned}\mathbb{E}[V_{\text{turn}} \mid \mathbf{G}_D] &= 3 \times S(2) \times q(4) \times p_{2,\text{surv}} = 3.66 \text{ VP/turn.} \\ \mathbb{E}[V_{\text{turn}} \mid \mathbf{G}_C] &= [S(4)p_{4,\text{surv}} + S(2)p_{2,\text{surv}}] \times q(3) \\ &= 2.43 \text{ VP/turn.}\end{aligned}$$

Notably, diversification yields a 50% higher rate. \square

These results suggest a natural build sequence. In the early game, the Cult should prioritize reaching $(2, 2, 2)$, building one suit at a time: $(2, 0, 0) \rightarrow (2, 2, 0) \rightarrow (2, 2, 2)$, since each new suit at $g = 2$ unlocks a draw bonus and scoring. Further gardens serve as attrition buffers but should not delay scoring: the delay value analysis (Section 4) shows $\Delta(2) < 0$, indicating that the Cult should score any suit at $g \geq 2$ immediately until the endgame adjustment (Section 4.3) overrides the base threshold.

6 MONTE CARLO SIMULATION

We simulate $N = 10,000$ games per configuration via an algorithm (Appendix D, Appendix E).

Table 1: Victory time by threshold (mean), $\mu = 0.4$.

Policy	$\lambda = 0.15$		$\lambda = 0.30$	
	Mean	95% CI	Mean	95% CI
Threshold-2	10.2	± 0.04	11.1	± 0.05
Threshold-3	11.8	± 0.04	13.0	± 0.06
Threshold-4	13.3	± 0.05	14.8	± 0.08

Table 2: Diversification comparison ($\mu = 0.4$, $\lambda = 0.2$, Diversified = $(2, 2, 2)$, Concentrated = $(4, 2, 2)$).

Strategy	Mean T	VP/turn	$\mathbb{P}(\text{win by } t=10)$
Diversified	10.0	3.1	0.67
Concentrated	10.9	2.9	0.48

In the simulation, the Concentrated strategy produces better results than expected (though still worse than Diversified). Inspecting trajectories shows that this is due

to the two strategies often being 'effectively' equivalent due to the difficulty of maintaining the $(2, 2, 2)$ position that allows constant full scoring.

Table 3: Sensitivity of $\mathbb{E}[T]$ to parameters ξ (base $\mu = 0.4$, $\lambda = 0.2$, $\theta = 2$).

Parameter	$\partial \mathbb{E}[T] / \partial \xi$	Interpretation
Attrition λ	+5.6	0.1 increase adds 0.56 turns
Growth μ	-17.5	0.1 increase saves 1.75 turns
Threshold θ	+1.7	1 unit adds 1.7 turns

Growth rate is the dominant parameter, and the hardest to model, given that it depends so heavily on a player's luck with Building and Conspiracies. In real play, higher μ is likely to also yield higher λ , as the higher growth from Sanctify conspiracies is only possible if the Cult's warriors are frequently attacked. The distribution is right-skewed with occasional long games ($T > 15$) due to poor draws or heavy attrition.

7 CONCLUSION

By developing a Markov chain framework for Lizard Cult scoring in *Root*, we found that $\theta^* = 2$ is robustly optimal because it is simultaneously the minimum viable scoring level and the draw bonus activation threshold. We also saw that the one-turn delay value $\Delta(g)$ is positive at $g = 1$ and negative at $g = 2$, showing that there is risk and no upside to waiting after obtaining 2 gardens in a suit. Ultimately, multi-suit scoring with configuration $(2, 2, 2)$ yields higher VP rates than blind growth.

REFERENCES

- [1] C. Wehrle, *Root: A Game of Woodland Might and Right & The Underworld Expansion*, Leder Games, 2019.
- [2] S. M. Ross, *Introduction to Probability Models*, 12th ed., Academic Press, 2019.
- [3] J. G. Kemeny and J. L. Snell, *Finite Markov Chains*, Springer-Verlag, New York, 1976.
- [4] N. L. Johnson, A. W. Kemp, and S. Kotz, *Univariate Discrete Distributions*, 3rd ed., Wiley, 2005.

APPENDIX

Full Appendix (A-F) is accessible through *The Delta Epsilon* online version.

A SPECIAL CASE OF TELESCOPIC SERIES

Sebastien Renard

This paper explores the convergent telescopic series $\sum_{n=1}^{\infty} \frac{1}{n(n+1)\dots(n+m)}$, where m is a nonzero positive integer. It starts by introducing the telescopic series and investigates its values for $m = 1, 2, 3, 4, 5, 6$, using a Python Program for $m = 4, 5, 6$, in order to formulate a conjecture, and proves in the last part that $\sum_{n=1}^{\infty} \frac{1}{n(n+1)\dots(n+m)} = \frac{1}{m \cdot m!}$. Although this report was written in January 2026, the investigation dates from June 2021: the author, being in high school, could not publish the paper at the time.

1 INTRODUCTION

Telescopic series are a type of series with a special property. Let $(a_n)_{n \in \mathbb{N}}$ be a sequence of real numbers with limit $l \in \mathbb{R} \cup \pm\infty$. The partial sum of order $n \in \mathbb{N}^*$ of the telescopic series associated to a_n is

$$S_n = \sum_{k=1}^n (a_k - a_{k-1}). \quad (1)$$

It can be rewritten as

$$S_n = \sum_{k=1}^n a_k - \sum_{k=1}^n a_{k-1} = \sum_{k=1}^n a_k - \sum_{k=0}^{n-1} a_k = a_n - a_0. \quad (2)$$

The equality (2) gives a direct value of the partial sum of a telescopic series, and, by taking the limit on both sides,

$$S = \lim_{n \rightarrow \infty} S_n = \lim_{n \rightarrow \infty} (a_n - a_0) = l - a_0. \quad (3)$$

For $m \in \mathbb{N}^*$, define the series $S^{(m)} = \sum_{n \in \mathbb{N}^*} \frac{1}{n(n+1)\dots(n+m)} = \sum_{n \in \mathbb{N}^*} \frac{1}{\prod_{i=0}^m (n+i)}$, is convergent as its general term is, $\forall n \in \mathbb{N}$, positive and bounded up by $\frac{1}{n^2}$ is convergent as its general term is, for all $n \in \mathbb{N}^*$, positive and bounded up by $\frac{1}{n^2}$ which is the general term of a convergent Riemann series. The nature and the value of the series for $m = 1$ and $m = 2$ are often presented as basic examples of convergent series in classes of calculus. In these examples, for $n \in \mathbb{N}^*$, the general terms $\frac{1}{n(n+1)}$ and $\frac{1}{n(n+1)(n+2)}$ can be decomposed using the partial fraction decomposition, resulting in the situation of (2), thus giving the value of the series by (3).

In the first part of this work, several values of the series $S^{(m)} = \sum_{n=1}^{\infty} \frac{1}{\prod_{i=0}^m (n+i)}$ for different values of $m \in \mathbb{N}^*$ will be given in order to formulate a conjecture with Python, which the second part of the work will focus on proving.

2 A CONJECTURE

To formulate a conjecture on the value of $S^{(m)}$ for all $m \in \mathbb{N}^*$, the series is computed for $m = 1$ using the partial fraction decomposition.

Let $k \in \mathbb{N}^*$. The roots of the polynomials X and $X + 1, 0$ and -1 , are distinct, so by the theorem on partial fraction decomposition, there exists $a, b \in \mathbb{R}$ such that $\frac{1}{k(k+1)} = \frac{a}{k} + \frac{b}{k+1}$. By multiplying both sides by k and plugging in $0, 1 = a + 0$ so $a = 1$. By multiplying both sides by $k + 1$ and plugging in $-1, -1 = b$. Then $\frac{1}{k(k+1)} = \frac{1}{k} - \frac{1}{k+1}$. By writing $a_k = -\frac{1}{k}$ for $k \in \mathbb{N}^*$ and by changing variables $k' = k + 1$, for $n \in \mathbb{N}^*$, $S_n = \sum_{k=2}^{n+1} (a_k - a_{k-1})$. $\lim_{n \rightarrow \infty} a_n = 0$. Then, by (3), $S^{(1)} = 0 - (-1) = 1$.

$$\sum_{n=1}^{\infty} \frac{1}{n(n+1)} = 1 = \frac{1}{1 \cdot 1!} \quad (4)$$

Now, for $m = 2$, for $k \in \mathbb{N}^*$, as $X, X + 1$ and $X + 2$ have distinct roots, so there exists $a, b, c \in \mathbb{R}$ such that $\frac{1}{k(k+1)(k+2)} = \frac{a}{k} + \frac{b}{k+1} + \frac{c}{k+2}$. By the previous method, $a = \frac{1}{2}, b = -1, c = \frac{1}{2}$. By writing $a_k = -\frac{1}{k}$ for $k \in \mathbb{N}^*$, for $n \in \mathbb{N}^*$, $S_n = -\frac{1}{2} \sum_{k=1}^n a_k + \sum_{k=1}^n a_{k+1} - \frac{1}{2} \sum_{k=1}^n a_{k+2}$, then $S_n = -\frac{1}{2} a_1 + \frac{1}{2} a_{n+1} + \frac{1}{2} a_2 - \frac{1}{2} a_{n+2}$. $\lim_{n \rightarrow \infty} a_{n+1} = \lim_{n \rightarrow \infty} a_{n+2} = 0$. Then

$$S^{(1)} = -\frac{1}{2} a_1 + \frac{1}{2} a_2 = \frac{1}{2} * 1 - \frac{1}{2} * \frac{1}{2} = \frac{1}{4}.$$

$$\sum_{n=1}^{\infty} \frac{1}{n(n+1)(n+2)} = \frac{1}{4} = \frac{1}{2 \cdot 2!} \quad (5)$$

For $m = 3$, the same method gives, for $k \in \mathbb{N}^*$, for $n \in \mathbb{N}^*$, $\frac{1}{k(k+1)(k+2)(k+3)} = \frac{1}{6k} - \frac{1}{2(k+1)} + \frac{1}{2(k+2)} - \frac{1}{6(k+3)}$. This gives

$$\sum_{n=1}^{\infty} \frac{1}{n(n+1)(n+2)(n+3)} = \frac{1}{24} = \frac{1}{3 \cdot 3!} \quad (6)$$

For values of $m \in \mathbb{N}^*$ with $m \geq 4$, a numerical approximation of $S^{(m)}$ was performed with Python, by computing the sums by adding the 1,000 first terms. For $m = 4$, the program returned

$$S^{(4)} = 0.0104166666666419138 \approx \frac{1}{4 \cdot 4!} = \frac{1}{96} \quad (7)$$

For $m = 5$, it returned

$$S^{(5)} = 0.0016666666666664705 \approx \frac{1}{5 \cdot 5!} = \frac{1}{600} \quad (8)$$

For $m = 6$, it returned

$$S^{(6)} = 0.00023148148148148003 \approx \frac{1}{6 \cdot 6!} = \frac{1}{4320} \quad (9)$$

These results lead to the following conjecture:

$$\sum_{n=1}^{\infty} \frac{1}{n(n+1)\dots(n+m)} = \frac{1}{m * m!}, \forall m \in \mathbb{N}^* \quad (10)$$

3 A THEOREM

Theorem 1. Let $m \in \mathbb{N}^*$. Then

$$\sum_{n=1}^{\infty} \frac{1}{n(n+1)\dots(n+m)} = \frac{1}{m * m!} \quad (11)$$

Proof. Let $m \in \mathbb{N}^* \setminus \{1\}$ and $k \in \mathbb{N}^*$. Let

$$u_k = \frac{1}{\prod_{i=0}^{m-1} (k+i)}. \quad (12)$$

Then,

$$\begin{aligned} u_{k+1} - u_k &= \frac{1}{\prod_{i=0}^{m-1} (k+1+i)} - \frac{1}{\prod_{i=0}^{m-1} (k+i)} \\ &= \frac{1}{\prod_{i=1}^m (k+i)} - \frac{1}{\prod_{i=0}^{m-1} (k+i)} \\ &= \frac{k - (k+m)}{\prod_{i=0}^m (k+i)} \\ &= \frac{-m}{\prod_{i=0}^m (k+i)} \end{aligned}$$

Then, isolating the inverse of the product,

$$\frac{1}{\prod_{i=0}^m (k+i)} = -\frac{1}{m} (u_{k+1} - u_k) \quad (13)$$

By summing on both sides over k from 1 to $n \in \mathbb{N}^*$,

$$\sum_{k=1}^n \frac{1}{\prod_{i=0}^m (k+i)} = -\frac{1}{m} (u_{n+1} - u_1) \quad (14)$$

As $\lim_{n \rightarrow \infty} u_n = 0$, taking the limits on both sides in (13) gives

$$\begin{aligned} S_m &= \sum_{k=1}^{\infty} \frac{1}{\prod_{i=0}^m (k+i)} = 0 - \left(-\frac{1}{m} u_1 \right) \\ &= \frac{1}{m} \cdot \frac{1}{\prod_{i=0}^{m-1} (1+i)} = \frac{1}{m} \cdot \frac{1}{\prod_{i=1}^m i} \\ &= \frac{1}{m * m!} \end{aligned}$$

Finally,

$$\sum_{n=1}^{\infty} \frac{1}{n(n+1)\dots(n+m)} = \frac{1}{m * m!} \quad (15)$$

□

4 ACKNOWLEDGEMENT

This work is entirely dedicated to Prof. Yves Moreau, the author's math teacher in high school. The author expresses him his deepest gratitude for helping him with the proof, as well as for changing his life trajectory and inspiring him to pursue studies in mathematics. The author also wants to thank his friend of many years Mr. Alexandre Vernet, who taught him the partial fraction decomposition method instead of listening to Usbek and Rica.

5 APPENDIX

The following Python program was used to approximate the values of $S(4), S(5), S(6)$ in order to conjecture the formula.

```
def f(m):
    s=0
    for k in range(1,1001):
        u=1
        for i in range(0,m+1):
            u=u*(k+i)
        s=s+1/u
    return s
```


theory of electromagnetism proposed by Maxwell due to the complicated nature of his equations.

2.2.1 A Brief Overview of the Original Maxwell's equations

We all know of the Maxwell's popular 4 equations. However, the original papers and treatise where his equations are cited were not just 4, but 20 equations, as it was written in Cartesian coordinates:

$$\begin{aligned}
 e + \frac{df}{dx} + \frac{dg}{dy} + \frac{dh}{dz} &= 0 \\
 \mu\alpha &= \frac{dH}{dy} - \frac{dG}{dz} \\
 \mu\beta &= \frac{dF}{dz} - \frac{dH}{dx} \\
 \mu\gamma &= \frac{dG}{dx} - \frac{dF}{dy} \\
 P &= \mu \left(\gamma \frac{dy}{dt} - \beta \frac{dz}{dt} \right) - \frac{dF}{dt} - \frac{d\Psi}{dx} \\
 Q &= \mu \left(\alpha \frac{dz}{dt} - \gamma \frac{dx}{dt} \right) - \frac{dG}{dt} - \frac{d\Psi}{dy} \\
 R &= \mu \left(\beta \frac{dx}{dt} - \alpha \frac{dy}{dt} \right) - \frac{dH}{dt} - \frac{d\Psi}{dz} \\
 \frac{d\gamma}{dy} - \frac{d\beta}{dz} &= 4\pi p' \quad p' = p + \frac{df}{dt} \\
 \frac{d\alpha}{dz} - \frac{d\gamma}{dx} &= 4\pi q' \quad q' = q + \frac{dg}{dt} \\
 \frac{d\beta}{dx} - \frac{d\alpha}{dy} &= 4\pi r' \quad r' = r + \frac{dh}{dt} \\
 P &= -\xi p \quad Q = -\xi q \quad R = -\xi r \\
 P &= kf \quad Q = kg \quad R = kh \\
 \frac{de}{dt} + \frac{dp}{dx} + \frac{dq}{dy} + \frac{dr}{dz} &= 0
 \end{aligned}$$

There had been many attempts to reformulate these equations but all were unsuccessful. One potential attempt was to incorporate **Quaternion** (\mathbb{H})³, which is the extension of the complex number system, and is mathematically represented as

$$a + b\mathbf{i} + c\mathbf{j} + d\mathbf{k}$$

where $\mathbf{i}^2 = \mathbf{j}^2 = \mathbf{k}^2 = \mathbf{ijk} = -1$.

To Heaviside, the usage of quaternion was unnatural, so he decided to stick with simple vectors and scalars. Furthermore, independently alongside Josiah Willard Gibbs, Heaviside introduce the notation of the curl and divergence operators, which are still used nowadays in vector calculus: $\nabla \times$ and $\nabla \cdot$.

³Originally, it was discovered by the Irish mathematician William Rowan Hamilton in 1843

2.2.2 Rewriting the Maxwell's Equations

His first moment of epiphany was during the summer of 1884, where he was able to relate the flow of energy (\mathbf{S}) in space with the electric (\mathbf{E}) and magnetic field (\mathbf{H})

$$\mathbf{S} = \mathbf{E} \times \mathbf{H}$$

With this, he was able to swiftly dismantle the Maxwell's equations one by one.

To heavily simplify what he has done to the 20 monstrosity of equations: he first removed the vector and scalars of potential \mathbf{A} and Ψ , respectively. In place, for the electromagnetic relationship, he used \mathbf{E} and \mathbf{H} (or, more fitting, the current convention \mathbf{B}) as terms of electric and magnetic field, respectively.

With Gauss' law, the entire summation of differential equation of the electric effects is reduced down to the divergence of the electric field. Similarly, Gauss' law for magnetism is simplified to the divergence of the magnetic field

$$\begin{aligned}
 \nabla \cdot \mathbf{E} &= \frac{\rho}{\epsilon_0} \\
 \nabla \cdot \mathbf{B} &= 0
 \end{aligned}$$

Next, most importantly, was relating the curl of the electric field ($\nabla \times \mathbf{E}$) to the rate of change of the magnetic field ($\frac{\partial \mathbf{B}}{\partial t}$) which enabled him to obtain Faraday's law

$$\nabla \times \mathbf{E} = -\frac{\partial \mathbf{B}}{\partial t}$$

Finally, relating the electric current (\mathbf{J}) and its production of the magnetic field to obtain the Ampere-Maxwell law

$$\nabla \times \mathbf{B} = \mu_0 \left(\mathbf{J} + \epsilon_0 \frac{\partial \mathbf{E}}{\partial t} \right)$$

And, at long last, the 4 Maxwell equations had been re-formalized. These continue to govern the entire study of classical electromagnetism to this day.

$$\begin{aligned}
 \nabla \cdot \mathbf{E} &= \frac{\rho}{\epsilon_0} \\
 \nabla \cdot \mathbf{B} &= 0 \\
 \nabla \times \mathbf{E} &= -\frac{\partial \mathbf{B}}{\partial t} \\
 \nabla \times \mathbf{B} &= \mu_0 \left(\mathbf{J} + \epsilon_0 \frac{\partial \mathbf{E}}{\partial t} \right)
 \end{aligned}$$

where \mathbf{E} is the electric field, \mathbf{B} is the magnetic field, \mathbf{J} is the current density, μ_0 for vacuum permeability and ϵ_0 for vacuum permittivity.

It must be noted that Heaviside did not present all 4 of these equations at once; rather, they were broken up into slow, incremental publications in the journal *The Electrician* during the 1880s.

2.3 Works on Inductive Loading

In 1886, Arthur Heaviside was experimenting with parallel circuit in telephone lines. He found that by adding more telephones to a single circuit, it improved the clarity of transmission. Turning to his brother Oliver for an explanation, Heaviside found that the leakage through each telephone line reduced the transmission distortion. Furthermore, increasing the circuit inductance (a term coined by Heaviside himself)⁴ by adding an inductor (typically a coil of wire), will significantly reduce the distortion. He essentially created an inductive load⁵. Regardless of this invention, Heaviside never patented his idea, which later in 1899 was taken by Michael Pupin, who sold the patent to the American Telephone & Telegraph (AT&T) company.

During the time of describing inductive loading, both Heaviside and his brother were prepared to publish their result to the *Journal of the Society of Telegraph Engineers and of Electricians*. However, it was promptly stopped by Arthur's superior, William Henry Preece. This was partly due to the fact that inductive loading was still part of their company project to fix transmission distortion. This work was even blocked by Preece when Heaviside tried to communicate it to *The Electricians*.

Luckily for the both of them, in November 1887, with the help of Lord Kelvin, Heaviside was able to persuade the journal *Philosophical Magazine* to accept this work alongside his series of articles on electromagnetic waves. At this point, Heaviside was slowly gaining attention in the scientific community. By 1891, he was elected a fellow of the Royal Society of London. In fact, his work made its way into many textbooks, most prominently, "Introduction to Maxwell's Theory of Electricity", which Albert Einstein and others would learn electromagnetism from.

3 LATER LIFE

In 1889, Heaviside and his parents moved to Paignton in Devonshire. Here, he continued his work in electromagnetic theory, but even more, the development of vector analysis and the usage of "operators" to solve differential equations. Eventually, this would be known as "operational calculus", which is a less rig-

⁴Inductance is a property of an electrical conductor, specifically the inductor, to oppose the change in current flow by storing the energy in a magnetic field

⁵A device that use the magnetic field from inductance to operate. Note that this is not the first instance where an inductive load is created, but it was Heaviside who helped bring it to real-world use.

orous technique in solving differential equations, but closely resembles our modern day method of Laplace transform. Unfortunately, the paper "On Operators in Physical Mathematics", where he outlined the usage of operational calculus, was rejected from the Royal Society proceedings by many pure mathematicians due to its lack of rigorous proof as well as his handling of divergent series.

In 1894, his mother passed away, which was followed by his father in 1896. That year, many of his friends at the Royal Society wanted to set up a grant from the Royal Society Relief Fund, but this was rejected by Heaviside as he did not want charity. However, later under the persuasion of his friends, the government offered him a pension of 120£/year (equivalent of 20,245.77£/year in 2026), which he accepted. This was later raised to 220£/year (equivalent of 37,117.24£/year in 2026).

Later in life, he suffered from paranoia and completely halted his scientific work in 1905. In 1908, his bother Charles arranged for him to move to Torquay, where he would share a large house and help with the mortgage payment with Mary Way, the sister of Charles' wife. His health improved, but due to his poor treatment of Mary Way, she promptly moved out in 1916, leaving him in complete isolation in his final years.

3.1 Death

Oliver Heaviside passed away on February 3rd, 1925 from a cerebral hemorrhage. He was buried on the same plot as his parents, where his name was added to the tombstone.

4 CONTRIBUTIONS

Although he did not have a massive impact on the scientific world, like many famous scientists of his time, he still had a large contribution to the community, some of which are still used to this day.

He introduced many electrical terms that are still widely used: impedance, conductance, inductance, permeability, permittance (now capacitance). Aside from just reformulating Maxwell's equations, he also pushed for the use of worked on vector calculus. He invented the Heaviside step function, which was used to model the sudden changes in electric circuit.

$$H(t) = \begin{cases} 0 & t < 0 \\ 1 & t \geq 0 \end{cases}$$

REFERENCES

- [1] Paul J. Nahin, *Oliver Heaviside: The Life, Work, and Times of an Electrical Genius of the Victorian Age* (Johns Hopkins University Press, 2002).
- [2] Hunt, B. J. Oliver Heaviside: A first-rate oddity. *Phys. Today* 65(11), 48 (2012); doi: 10.1063/PT.3.1788. <http://dx.doi.org/10.1063/PT.3.1788>.
- [3] Britannica Editors. "Oliver Heaviside." *Encyclopedia Britannica*, January 30, 2026. <https://www.britannica.com/biography/Oliver-Heaviside>.
- [4] Maxwell, James Clerk (1865). "A dynamical theory of the electromagnetic field". *Philosophical Transactions of the Royal Society of London*. 155: 459–512. doi:10.1098/rstl.1865.0008. OL 25533062M. S2CID 186207827
- [5] M. R. Spiegel, *Applied Differential Equations*, 3rd ed. Pearson, 1980. ISBN: 978-0-13-040097-0.

JOKES

What is one thing stochastic processes like to do?

– They like to take a Random Walk

FERMAT'S LAST CROSSWORD

Max Gross

1	2	3	4		5	6	7	8	9		10	11	12	13
14					15						16			
17					18					19				
	20			21										
			22				23					24	25	26
27	28	29		30		31			32		33			
34			35			36	37		38					
39					40				41		42			
43				44		45				46				
47					48			49				50		
51				52		53	54		55		56			
		57	58					59				61	61	
62	62											64		65
66					67							68		
69					70							71		

ACROSS

- 1 (In) place
- 5 Finishes running, as a Turing Machine
- 10 Diminutive suffix
- 14 Raring
- 15 Greek theatre
- 16 "Is there a question?"

DOWN

- 1 Velázquez's "___ Meninas"
- 2 "Indubitably"
- 3 Homeland for William Rowan Hamilton
- 4 Upstate New York town
- 5 "Such a unique solution"
- 6 Out of focus?

- 17 Tech support?
- 18 “That lecture was incredible!”
- 20 Course Assistants via Teaching Assistants, often
- 22 Donnybrook
- 23 ___ -à- ___
- 24 Ballpark figure?
- 27 Abel Prize winner Widgerson
- 30 (Former) Bolivian president Morales
- 32 McGill grad Trudeau
- 34 Limit of a series, perhaps
- 38 Mathematics educator Emma Castelnuovo vis-à-vis algebraic geometer Federigo Enriques
- 39 Ntwrk. covering the Olympics, but not the Olympiad
- 40 Sleds supine
- 42 Swiss luxury watch manufacturer
- 43 “The Prince of Mathematicians”
- 45 The Principles of Quantum Mechanics by Paul Dirac, e.g.
- 47 Late payment
- 49 El Niño component
- 50 Grassy knoll
- 51 March 14th treat
- 52 Rapper who coined the slang term ‘based’
- 55 Double helix material
- 57 Integral field of study... or a description for 20-Across, 34-Across, and 45-Across
- 62 Quality of machine learning or some statistical modelling
- 64 Looked over
- 66 Snog in Seville
- 67 Savouriness
- 68 “Once upon a time” story
- 69 Grp.
- 70 ___ Puente (archaeological sight of the Maya)
- 71 Baldwin brother
- 7 Infimum, to normal people
- 8 Canadian code breaker and mathematician William
- 9 Acumen
- 10 Merit
- 11 Buster Keaton film taking place in the American Civil War
- 12 London Fog, for instance
- 13 Ambulance VIPs
- 19 Szechuan Sauce for Mulan, for one
- 21 MATH 315 subjs.
- 25 Instruction for a course?
- 26 Black Canadian businesswoman Curry
- 27 May-December romance feature
- 28 “Lives of the Most Excellent Painters, Sculptors, and Architects” author
- 29 Quality of Applied Mathematics, pejoratively
- 31 Egg cell
- 33 French fashion house
- 35 Start of Massachusetts’ state motto
- 36 “Me” problem
- 37 Hawaiian goose
- 41 West Side Story lyricist (and onetime math major) Stephen
- 44 Pheasant ragout
- 46 Component of 55-Across
- 48 Baby Back, for example
- 53 “I can do it”
- 54 Stanford professor of 57-Across Conrey
- 56 Major vessel
- 58 Japanese noodle
- 59 “This program is specifically designed to be viewed by adults...” rating
- 60 Not complex?
- 61 Institution where Benoit Mandelbrot was the oldest professor in their history to receive tenure
- 62 Org. for keglers
- 63 New, e.g. Abbr.
- 65 Exams mo.

CHEEGER'S INEQUALITY

Olivia Choi

In this report, we present Cheeger's inequality in three sections. The first introduces a motivating example and discusses the relevant combinatorial notions, including graph cuts and conductance. The second examines the spectral measure of a graph, focusing on the Rayleigh quotient, normalized Laplacian, and Fiedler value. The third states Cheeger's inequality along with its proofs, highlighting its significance in linking a combinatorial and a spectral concept.

Before we officially begin, let us recall the definition of a graph. A graph $G = (V, E)$ consists of a finite set of vertices V , and a set of edges $E \subseteq \{\{u, v\} \mid u, v \in V\}$, where each edge connects a pair of vertices. Throughout this report, we assume that G is a *simple, undirected* graph (unless otherwise specified).

Definition. (Simple Graph) A *simple graph* is an undirected graph that contains no graph loops or multiple edges. A simple graph may be connected or disconnected.

Definition. (Undirected and Directed Graphs) An *undirected graph* is a graph where the edges do not have a specific direction and is bidirectional in nature.

A *directed graph* is a unidirectional graph, where the edges have a specific direction. Additionally, for a directed edge uv , we call the starting vertex u the *tail*, and the ending vertex v the *head*.

1 GRAPH PARTITIONING AND CONDUCTANCE

In this section, we introduce graph partitioning and conductance with a motivating example and then define them mathematically.

1.1 A Motivating Example

Consider that we want to partition a graph. This is called a *cut*.

Definition. (Cut) A *cut* in a graph is a partition of the vertex set V into two disjoint subsets S and $V \setminus S$.

Think of a graph with six vertices with one bottleneck edge (as in Figure 1). This graph can be partitioned in several ways (the cut in Figure 1 shows the type of "good" partition we will discuss shortly).




Figure 1: A simple graph with six vertices with one bottleneck edge.

This naturally raises the question: what constitutes a "good" partition? Informally, a "good" partition divides the graph into two parts such that there are few edges between the groups and many edges within each group. Notice in our example some partitions separate densely connected clusters with few cross edges. Our goal, of course, is to quantify how good a partition is.

Let $S \subseteq V$, and $\bar{S} = V \setminus S$. Then, $|\partial S|$ = number of edges with one endpoint in S and the other in \bar{S} .

But this alone is not enough. A small $|\partial S|$ is meaningless if S is tiny. So, we normalize by the "volume" of S or \bar{S} , which gives us *conductance*. The conductance will help us accurately quantify how "good" a cut is.

1.2 Conductance and Cheeger Constant

Definition. (Degree) The *degree* of a vertex, denoted $\deg(v)$, is the number of edges that are connected to that vertex.

Definition. (Volume) The *volume* of a subset $S \subseteq V$ measures how "large" S is in terms of the total degree of its vertices

$$\text{vol}(S) = \sum_{v \in S} \deg(v).$$

Definition. (Conductance) A *conductance* of a cut (S, \bar{S}) is defined as

$$\phi(S) = \frac{|\partial S|}{\min\{\text{vol}(S), \text{vol}(\bar{S})\}}. \quad (1)$$

A *good cut* minimizes the number of crossing edges while maintaining roughly balanced volume. So, a "good" cut has low conductance.

Definition. (Cheeger Constant) The *Cheeger constant* is defined as

$$h(G) = \min_{S \subseteq V} \phi(S), \text{ where } 0 < \text{vol}(S) \leq \frac{1}{2} \text{vol}(V). \quad (2)$$

The restriction to subsets with volume at most half ensures that we avoid trivial cuts involving very small

sets. The Cheeger constant $h(G)$ captures the minimal conductance across all non-trivial bipartitions.

So, each subset S gives one value of $\phi(S)$, and the Cheeger constant is the minimum of these values over all nontrivial subsets. Returning to our example (1.1), since the bottleneck edge is almost disconnected, we predict that the Cheeger constant is close to zero.

2 RAYLEIGH QUOTIENT AND NORMALIZED LAPLACIAN

In this section, we introduce the Rayleigh quotient, explain its connection to the eigenvalues of the normalized Laplacian, and show how it relates spectral quantities to graph structure, particularly through the Laplacian matrix.

2.1 Key Definitions

Before discussing the Rayleigh quotient, we need some definitions:

Definition. (Adjacency Matrix) The *adjacency matrix* $A = (A_{ij})$ is defined as

$$A_{ij} = \begin{cases} 1, & \text{if } i \sim j \\ 0, & \text{otherwise,} \end{cases} \quad (3)$$

where $i \sim j$ means the vertices i and j are connected by an edge in G .

Remark. We shall mostly discuss undirected graphs, that is, if there is an edge from i to j , then there is an edge from j to i (We have also previously defined undirected graphs in section 1). For such graphs, the adjacency matrix A_{ij} is symmetric. Additionally, for graphs without loops, the diagonal entries of A are zero.

Definition. (Path) A *path* is a finite or infinite sequence of edges that joins a sequence of vertices.

Definition. (Degree Matrix) The *degree matrix* D is a diagonal matrix with each diagonal entry equal to the degree of that vertex (recall that we have defined degree previously in section 1):

$$D_{ii} = \text{deg}(i) \quad (4)$$

2.2 The Rayleigh Quotient

Definition. Let $B \in \mathbb{R}^{n \times n}$ be a real symmetric matrix (This matrix is typically the normalized Laplacian, which we later discuss in (8)). For a nonzero vector $x \in \mathbb{R}^n \setminus \{0\}$, the *Rayleigh quotient* with respect to B is:

$$R_B(x) = \frac{x^T B x}{x^T x} \quad (5)$$

This quantity is always real, and it describes how the quadratic form induced by B scales relative to the squared length of x . For symmetric matrices, the Rayleigh quotient plays a key role in characterizing eigenvalues and eigenvectors.

Lemma 1. If x is an eigenvector of B with eigenvalue λ , then $R_B(x) = \lambda$. Thus, eigenvalues are stationary values of the Rayleigh quotient.

To see this, first note that the stationary points of $R_B(x)$ have a zero gradient. Using the quotient rule and the fact that B is symmetric,

$$\nabla R_B(x) = \frac{2Bx(x^T x) - 2x(x^T Bx)}{(x^T x)^2} = 0.$$

By multiplying both sides with $(x^T x)^2$ and rearranging gives us $Bx = \frac{x^T Bx}{x^T x} x = R_B(x)x$, which is exactly the definition of an eigenvector equation with the eigenvalue $\lambda = R_B(x)$.

This observation motivates the use of the Rayleigh quotient to estimate or bound eigenvalues. The Rayleigh Quotient is a variational tool. Our goal is to define a matrix B such that minimizing $R_B(x)$ gives us a value related to $h(G)$.

2.3 The Laplacian as a Quadratic Form

Let $G = (V, E)$ be a simple, undirected, connected graph. Let A denote the *adjacency matrix* (3) and D the *degree matrix* (4) of G .

Definition. (Unnormalized Laplacian) The *unnormalized Laplacian* is defined as

$$L = D - A. \quad (6)$$

For the unnormalized Laplacian L of $G = (V, E)$, the associated quadratic form is

$$x^T L x = \sum_{i \sim j} (x_i - x_j)^2. \quad (7)$$

Proof. First, observe that

$$x^T D x = \sum_{i \in V} d_i x_i^2,$$

and since A is symmetric for an undirected graph, the term $x^T A x$ counts the product $x_i x_j$ twice for every edge $\{i, j\}$,

$$x^T A x = \sum_{i \sim j (\text{ordered})} x_i x_j = 2 \sum_{\{i, j\} \in E} x_i x_j.$$

Thus, the quadratic form of the unnormalized Laplacian is:

$$\begin{aligned}
 x^T Lx &= x^T (D - A)x \\
 &= x^T Dx - x^T Ax \\
 &= \sum_{i \in V} d_i x_i^2 - 2 \sum_{\{i,j\} \in E} x_i x_j \\
 &= \sum_{\{i,j\} \in E} (x_i^2 + x_j^2) - 2 \sum_{\{i,j\} \in E} x_i x_j \\
 &= \sum_{\{i,j\} \in E} (x_i^2 - 2x_i x_j + x_j^2) \\
 &= \sum_{\{i,j\} \in E} (x_i - x_j)^2 = \sum_{i \sim j} (x_i - x_j)^2.
 \end{aligned}$$

□

This measures how “smooth” the vector x is over the graph. If neighboring vertices have similar values, the quadratic form is small; if they differ strongly, the value is large. In the context of the Rayleigh quotient, minimizing $R_L(x) = \frac{x^T Lx}{x^T x}$ seeks vectors x that are smooth relative to their squared length.

2.4 The Normalized Laplacian

Although the unnormalized quadratic form $x^T Lx$ measures total variation, it is inherently biased toward small sets, as a cut isolating a single high-degree vertex can lead to a large value. So like we did with $|\partial S|$ in Section 1.2, to account for the relative size (or volume) of partitions, we must normalize the Laplacian.

Definition. (Normalized Laplacian) *The normalized Laplacian is defined as*

$$\hat{L} = D^{-1/2} L D^{-1/2} = I - D^{-1/2} A D^{-1/2}. \quad (8)$$

These matrices are central in spectral graph theory because their eigenvalues encode structural properties of the graph, such as connectivity and expansion. In particular, the eigenvalues of \hat{L} interact naturally with the Rayleigh quotient to provide a volume-balanced measure of variation.

Through the Rayleigh quotient

$$R_{\hat{L}}(x) = \frac{x^T \hat{L} x}{x^T x},$$

we see that minimizing this quantity over nonzero vectors x leads directly to eigenvalues of \hat{L} . Thus, the Laplacian quadratic form not only captures variation across the graph but also provides the variational characterization of eigenvalues, linking combinatorial structure with spectral properties.

2.5 Understanding Eigenvalues of the Laplacian

We now examine what the Laplacian's structure implies about its eigenvalues.

The matrix has several important properties:

- L is symmetric and real-valued, so all its eigenvalues are real.
- L is positive semi-definite, meaning all its eigenvalues are non-negative.
- The smallest eigenvalue of L is always 0, and its eigenvector is any constant vector.

This last fact can be understood as follows: if we multiply the Laplacian matrix L by a constant vector, we get zero because each row sums to zero. That is,

$$L \cdot \mathbf{1} = 0$$

This shows that 0 is always an eigenvalue of L , and its eigenvector is any constant vector.

Now suppose $x \in \mathbb{R}^n$ is not a constant vector, that is, there exists i, j such that $x_i \neq x_j$. Then at least one term $(x_i - x_j)^2$ in the expression

$$x^T Lx = \sum_{(i,j) \in E} (x_i - x_j)^2 \quad (9)$$

will be strictly positive. Then the quadratic form $x^T Lx$ is strictly positive, so the Rayleigh quotient $R_L(x) = \frac{x^T Lx}{x^T x}$ is also strictly positive.

So, if the graph is connected, 0 is a simple eigenvalue and all other eigenvalues are strictly positive. The first non-zero eigenvalue is especially meaningful here, and we give its own definition.

Definition. (Fiedler value) *The Fiedler value, λ_2 , is the second-smallest (first non-zero) eigenvalue of the normalized Laplacian \hat{L} and reflects how well-connected the graph is. A small λ_2 implies a bottleneck or sparse cut, and a large λ_2 suggests strong connectivity.*

We later see that this eigenvalue is precisely the spectral quantity that appears in Cheeger's inequality.

To give some insight to the readers, since the first eigenvector of \hat{L} is $D^{1/2} \mathbf{1}$, the Courant-Fischer Theorem states that the second smallest eigenvalue λ_2 can be found by minimizing the Rayleigh quotient over all vectors x that are orthogonal to $D^{1/2} \mathbf{1}$. In terms of the original vertex values, this is equivalent to vector g such that $\sum v g(v) d_v = 0$.

3 CHEEGER'S INEQUALITY

Cheeger's inequality connects the combinatorial isoperimetric properties of a graph (i.e., how "tightly" a subset of vertices is connected to the rest of the graph; via conductance) with its spectral properties (specifically, the second smallest eigenvalue of the normalized Laplacian, λ_2). It provides both a lower and an upper bound on the Cheeger constant in terms of the λ_2 .

Theorem 2. (Cheeger Inequality) *Let $G = (V, E)$ be a graph, λ_2 be the Fiedler value of \hat{L} , and $h(G)$ the Cheeger constant. The Cheeger inequality states that*

$$\frac{\lambda_2}{2} \leq h(G) \leq \sqrt{2\lambda_2}. \quad (10)$$

Before presenting the proofs, we informally interpret the inequality. The lower bound tells us that if there's a sparse cut, then λ_2 must be small. The upper bound tells us a small λ_2 guarantees the existence of a sparse cut.

This is what makes the inequality useful – a spectral quantity, λ_2 , can approximate a combinatorial problem: finding a near-optimal cut.

Remark. Cheeger's inequality has higher-order extensions that are useful when partitioning the graph into more than two subsets.

Definition. (Higher-Order Cheeger Constant) *When partitioning the graph into $k \geq 2$ disjoint nonempty subsets S_1, \dots, S_k , one defines the k -way Cheeger constant*

$$h_k(G) = \min_{\text{disjoint } S_1, \dots, S_k} \max_{1 \leq i \leq k} \phi(S_i),$$

where $\phi(S_i)$ is the conductance of S_i . This measures the quality of the *worst* of the k cuts.

Theorem 3. (Higher-Order Cheeger Inequality) *The higher-order Cheeger inequality then relates this to the k -th smallest eigenvalue λ_k of \hat{L} . In particular, it states that*

$$\frac{\lambda_k}{2} \leq h_k(G) \leq C \cdot k^2 \cdot \sqrt{\lambda_k},$$

for some constant C . This is due to works by Lee, Ghahramani, and Trevisan [1], and the restrictions to this inequality can also be found here.

Remark. The definition of $h_k(G)$ requires each subset in the k -partition to have size of at most $|V|/2$, ensuring non-trivial expansion (as in the classical case). Moreover, the bounds deteriorate with increasing k , and the inequality becomes more trivial if G is disconnected. A detailed discussion of these restrictions can be further found in [1].

3.1 Proof of the Lower Bound

Let us first prove the lower bound of Cheeger's inequality:

$$\frac{\lambda_2}{2} \leq h(G).$$

We follow the approach presented by Hu [2] in his proof for the lower bound.

Proof. We use the variational definition of λ_2 to do this.

Let $S \subseteq V$ be a subset of vertices such that

- $\text{vol}(S) \leq \frac{1}{2} \text{vol}(V)$, and
- $h(G) = \frac{|\partial S|}{\text{vol}(S)}$

We define a function $g : V \rightarrow \mathbb{R}$ as follows:

$$g(v) = \begin{cases} 1, & \text{if } v \in S \\ -\frac{\text{vol}(S)}{\text{vol}(\bar{S})}, & \text{if } v \in \bar{S}. \end{cases}$$

This function is chosen so that $g \perp \mathbf{1}$ with respect to the degree inner product

$$\sum_v g(v) d_v = 0.$$

We now compute the Rayleigh quotient for g with respect to \hat{L} :

$$R_{\hat{L}}(g) = \frac{g^T \hat{L} g}{g^T g}.$$

Recall from Section 2.3 (Equation 7) that for the unnormalized Laplacian, $x^T L x = \sum_{i \sim j} (x_i - x_j)^2$. By substituting the definition of the normalized Laplacian $\hat{L} = D^{-1/2} L D^{-1/2}$ and letting $x = D^{1/2} g$, we obtain the normalized identity

$$g^T \hat{L} g = \sum_{(i,j) \in E} \left(\frac{g(i)}{\sqrt{d_i}} - \frac{g(j)}{\sqrt{d_j}} \right)^2.$$

Now, we bound this. Since all edges crossing from S to \bar{S} contribute a nonzero term to the sum, and g is constant on S and \bar{S} , the only nonzero contributions come from these crossing edges.

We get

$$R_{\hat{L}}(g) \leq 2 \cdot \phi(S) = 2 \cdot h(G).$$

Additionally, we have defined λ_2 as the minimum Rayleigh quotient over all degree-orthogonal vectors. Thus,

$$\lambda_2 \leq R_{\hat{L}}(g) \leq 2h(G) \implies \frac{\lambda_2}{2} \leq h(G).$$

This completes the proof of the lower bound. \square

3.2 Proof of the Upper Bound

We now prove the upper bound of Cheeger's inequality:

$$h(G) \leq \sqrt{2\lambda_2}.$$

This section is motivated by the proofs in Chung [3]. In particular, we follow the proof using the eigenvector associated with λ_2 of the normalized Laplacian to construct a set whose conductance approximates the Cheeger constant.

Proof. Let f be a real-valued eigenvector satisfying

$$\hat{L}f = \lambda_2 f,$$

and normalized so that $f \perp \mathbf{1}$, meaning

$$\sum_{v \in V} f(v) \cdot d_v = 0.$$

This orthogonality condition ensures that f is not a constant vector and plays a key role in isolating the structure of the graph beyond its trivial connectedness.

We sort the vertices v_1, v_2, \dots, v_n so that

$$f(v_1) \geq f(v_2) \geq \dots \geq f(v_n).$$

For each index i , define the set

$$S_i = \{v_1, v_2, \dots, v_i\},$$

and for each S_i , compute the conductance (1)

$$\phi(S_i) = \frac{|\partial S_i|}{\min\{\text{vol}(S_i), \text{vol}(V \setminus S_i)\}}.$$

Let $\alpha = \min_i \phi(S_i)$. The idea here is that while there are exponentially many cuts, we only consider cuts defined by ordering vertices according to the values of f , leading to only $n - 1$ possible sets S_i . This is a significant reduction, and the Cheeger inequality guarantees that one of these sets will yield a conductance not much worse than the optimal.

The remaining part of the proof is algebraic. We define the auxiliary function $g(v) = f(v)/\sqrt{d_v}$ and consider only its positive part $g_+ = \max(0, g)$. From the variational properties of λ_2 , the Rayleigh quotient of g_+ is bounded by λ_2 .

Applying the Cauchy-Schwarz inequality to $|g_+^2(i) - g_+^2(j)|$, we can relate the squared differences in the Rayleigh quotient to the linear edge counts in the conductance.

Through a summation by parts (the sum-over-cuts identity), we show that the sum of these absolute differences is at least $\alpha \cdot \sum_v g_+^2(v) d_v$.

Finally, using the inequality

$$\sum_{\{i,j\} \in E} (g_+(i) + g_+(j))^2 \leq 2 \sum_{v \in V} g_+^2(v) d_v,$$

we combine these bounds to obtain that

$$\lambda_2 \geq \frac{\alpha^2}{2} \implies \alpha \leq \sqrt{2\lambda_2}.$$

Since the Cheeger constant $h(G)$ is the minimum conductance over all possible cuts, it must be that $h(G) \leq \alpha$. This completes the proof of the upper bound:

$$h(G) \leq \sqrt{2\lambda_2}. \quad \square$$

Therefore, with the proof of the lower bound in (3.1), we get the Cheeger's inequality we stated in Theorem 2:

$$\frac{\lambda_2}{2} \leq h(G) \leq \sqrt{2\lambda_2}.$$

REFERENCES

- [1] Lee, James; Gharan, Oveis; Trevisan, Luca. *Multi-way spectral partitioning and higher-order Cheeger inequalities*. In Proceedings of the 44th ACM Symposium on Theory of Computing (STOC), p. 1117 – 1130. 2012. Available at: <https://arxiv.org/abs/1111.1055>
- [2] Hu, William. *Graph partitioning and multi-way Cheeger inequalities*. University of Chicago REU, p. 8 – 10. 2023. Available at: <https://math.uchicago.edu/~may/REU2023/REUPapers/Hu,William.pdf>
- [3] Chung, Fan. *Four proofs for the Cheeger inequality and graph partition algorithms*. In Proceedings of ICCM, vol. 2, p. 378. 2007.

MULTIVARIATE COMPOSITE COPULAS: THEORY AND APPLICATION TO GERMAN GOVERNMENT BOND YIELDS

Gabriella Chen

In this paper, we investigate multivariate composite copulas as a flexible approach to modeling dependence structures. Unlike standard copula families, composite copulas generate new models through composition. We apply this framework to German government bond yields and show improved modeling of asymmetric tail dependence.

1 INTRODUCTION

1.1 The Copula Construction Problem

Since Sklar's fundamental theorem [2], copula theory has revolutionized multivariate dependency modelling. This theorem states that any multivariate distribution can be decomposed into its marginal distribution and a copula function capable of capturing dependence structures. While this theoretical foundation is robust, the practical challenge is constructing a family of copula functions that is both flexible enough to capture complex real-world dependencies and easy to estimate and simulate.

Traditional copula families, including elliptical copulas (Gaussian, t) and Archimedean copulas (Clayton, Gumbel, Frank), have well-documented limitations. Elliptical copulas impose radial symmetry, while Archimedean copulas are limited to exchangeable structures and specific tail dependence patterns. These limitations often render them insufficient to capture complex real-world data, such as the subtle dependence structures observed in financial markets, particularly asymmetric tail dependencies and different correlation mechanism.

1.2 Preliminary Concepts and Notation

An n -dimensional copula is a function

$$C : [0, 1]^n \rightarrow [0, 1], \quad (u_1, \dots, u_n) \in [0, 1]^n,$$

satisfying:

1. $C(u_1, \dots, u_n) = 0$ if any $u_i = 0$,
2. $C(1, \dots, 1, u_i, 1, \dots, 1) = u_i$ for all $i = 1, \dots, n$,
3. C is n -increasing, meaning that for any rectangle $\prod_{i=1}^n [a_i, b_i] \subset [0, 1]^n$, the C -volume

$$\sum_{\varepsilon_i \in \{0,1\}} (-1)^{n-\sum \varepsilon_i} C(u_1^{(\varepsilon_1)}, \dots, u_n^{(\varepsilon_n)}) \geq 0,$$

where $u_i^{(0)} = a_i$ and $u_i^{(1)} = b_i$. Intuitively, this condition ensures that C assigns nonnegative "probability mass" to every rectangular region, so it behaves like a joint distribution function.

Fundamental examples of copulas include:

- The *independence copula*

$$\Pi(u_1, \dots, u_n) = \prod_{i=1}^n u_i,$$

which corresponds to the case where all variables are independent.

- The *Fréchet–Hoeffding upper bound*

$$M(u_1, \dots, u_n) = \min\{u_1, \dots, u_n\},$$

which represents maximal positive dependence. Intuitively, this corresponds to the case where all variables move together (comonotonicity).

- The *Fréchet–Hoeffding lower bound*

$$W(u_1, \dots, u_n) = \max\{u_1 + \dots + u_n - n + 1, 0\},$$

which represents maximal negative dependence. This bound is only a copula in the case $n = 2$, where it corresponds to perfect negative dependence (countermonotonicity).

1.3 Function Classes and Generalized Inverses

The composite copula construction relies on carefully defined classes of bivariate functions. Let $\mathcal{F}_{1,R-I}$ denote the family of bivariate functions $f : [0, 1]^2 \rightarrow [0, 1]$ such that for each fixed $y \in [0, 1]$, the function $f(x, y)$ is right-continuous and increasing in x . Similarly, define $\mathcal{F}_{1,L-D}$ for left-continuous decreasing functions, and analogous classes for fixed x with respect to y .

For $f \in \mathcal{F}_{1,R-I}$ and fixed $y \in [0, 1]$, the generalized inverse is defined as:

$$f^{[-1]}(u|\cdot, y) = \inf\{x \in [0, 1] : f(x, y) \geq u\} \quad (1)$$

These generalized inverses satisfy key properties essential for the composite construction:

$$f^{[-1]}(u|\cdot, y) \leq x \Leftrightarrow u \leq f(x, y) \quad \text{for } f \in \mathcal{F}_{1,R-I} \quad (2)$$

$$x \leq f^{(-1)}(u|\cdot, y) \Leftrightarrow u \leq f(x, y) \quad \text{for } f \in \mathcal{F}_{1,L-D} \quad (3)$$

1.4 Report Structure

Section 2 introduces the complete mathematical framework of multivariate composite copula function. Section 3 gives the numerical illustration of multivariate copulas. Section 4 details the German government bond dataset and conducts some preliminary empirical dependency analysis. Section 5 explores the limitations of some commonly used copulas in this application. Section 6 discusses the advantages and implementation consideration of the composite copula method.

2 MATHEMATICAL FOUNDATIONS OF MULTIVARIATE COMPOSITE COPULAS

2.1 Composite Copula Definition and Existence Theorem

The multivariate composite copula framework, introduced by [1], represents a new copula construction method. Rather than proposing another copula family, this framework provides a generative approach for creating new copulas through the composition of existing ones. The formal definition of the multivariate composite copula is as follows:

Definition (Multivariate Composite Copula). *Given two n -dimensional copulas B and C , and a vector of bivariate functions $\mathbf{f} = (f_1, \dots, f_n)$ with $f_i : [0, 1]^2 \rightarrow [0, 1]$, the composite function $B \overset{\mathbf{f}}{\circ} C : [0, 1]^n \rightarrow [0, 1]$ is defined as:*

$$B \overset{\mathbf{f}}{\circ} C(u_1, \dots, u_n) = \mathbb{E}[B(f_1(u_1, U_1), \dots, f_n(u_n, U_n))] \quad (4)$$

where (U_1, \dots, U_n) is a random vector with distribution C , and $(u_1, \dots, u_n) \in [0, 1]^n$.

The fundamental theoretical result of [1] establishes when this construction yields a valid copula:

Theorem 1 (Existence and Characterization). *Suppose B and C are n -dimensional copulas and $\mathbf{f} = (f_1, \dots, f_n)$ satisfies:*

1. For each $i = 1, \dots, n$ and fixed $y \in [0, 1]$, $f_i(x, y)$ is increasing in x
2. For each $i = 1, \dots, n$ and $x \in [0, 1]$, $\int_0^1 f_i(x, y) dy = x$

Then $B \overset{\mathbf{f}}{\circ} C$ is an n -dimensional copula. Moreover, if $f_i \in \mathcal{F}_{1,R-I}$, then $f_i^{[-1]}(V_i | \cdot, U_i)$ is a Uniform $[0, 1]$ random variable and the random vector $(f_1^{[-1]}(V_1 | \cdot, U_1), \dots, f_n^{[-1]}(V_n | \cdot, U_n))$ follows the distribution $B \overset{\mathbf{f}}{\circ} C$, where $(V_1, \dots, V_n) \sim B$ is independent of $(U_1, \dots, U_n) \sim C$.

2.2 Constructing Valid Function Vectors

A general method for constructing valid function vectors \mathbf{f} utilizes the theory of modified partial Dini derivatives. For any bivariate copula D , define the modified partial Dini derivative with respect to the second variable as:

$$\mathcal{D}_2 D(x, y) = \begin{cases} \inf_{u>x} \mathcal{D}_2^+ D(u, y), & 0 \leq y < 1 \\ & 0 \leq x \leq 1 \\ \inf_{u>x} \mathcal{D}_2^- D(u, 1), & y = 1 \\ & 0 \leq x \leq 1 \end{cases} \quad (5)$$

where \mathcal{D}_2^+ and \mathcal{D}_2^- denote the right and left derivatives, respectively.

Theorem 2 (Function Construction via Copulas). *For any bivariate copula D , the function $f(x, y) = \mathcal{D}_2 D(x, y)$ satisfies both assumptions of the existence Theorem 1. Moreover, if D is absolutely continuous, then $\mathcal{D}_2 D(x, y) = \frac{\partial}{\partial y} D(x, y)$ almost everywhere.*

This construction provides an intuitive interpretation: when the conditional distribution $\mathbb{P}(X \leq x | Y = y)$ is continuous in y , we can take $f(x, y) = \mathbb{P}(X \leq x | Y = y)$, which corresponds to $\mathcal{D}_2 D(x, y)$ for the copula D of (X, Y) .

2.3 Probabilistic Structure and Simulation Methodology

The composite copula admits a clear probabilistic interpretation through its stochastic representation. When the conditions of the theorem are satisfied, we have the equivalence:

$$(X_1, \dots, X_n) \sim B \overset{\mathbf{f}}{\circ} C \iff X_i = f_i^{[-1]}(V_i | \cdot, U_i), \quad (6)$$

$$i = 1, \dots, n$$

This representation leads to a straightforward simulation algorithm:

The effectiveness of this algorithm stems from the properties of the probability integral transformation and the generalized inverse relation.

2.4 Properties and Special Cases

The composite copula framework preserves many desirable properties, as shown by [1]:

Property 2 (Marginality). Suppose that $\mathbf{f} = (f_1, \dots, f_n) \in \mathcal{F}$, and B and C are n -dimensional

copulas. For each $i = 1, \dots, n$, denote $\mathbf{f}_{-i} = (f_1, \dots, f_{i-1}, f_{i+1}, \dots, f_n)$. Then:

$$B \overset{\mathbf{f}}{\circ} C(u_1, \dots, u_{i-1}, 1, u_{i+1}, \dots, u_n) = B_{-i} \overset{\mathbf{f}_{-i}}{\circ} C_{-i}(u_1, \dots, u_{i-1}, u_{i+1}, \dots, u_n) \quad (7)$$

where

$$B_{-i}(u_1, \dots, u_{i-1}, u_{i+1}, \dots, u_n) := B(u_1, \dots, u_{i-1}, 1, u_{i+1}, \dots, u_n) \quad (8)$$

and

$$C_{-i}(u_1, \dots, u_{i-1}, u_{i+1}, \dots, u_n) := C(u_1, \dots, u_{i-1}, 1, u_{i+1}, \dots, u_n) \quad (9)$$

are the $(n-1)$ -marginal copulas of B and C , respectively.

Property 2 (Reproduction). The framework reproduces important special cases:

$$B \overset{\mathbf{f}}{\circ} C = B, \text{ when } f_i(x, y) = \mathcal{D}_2 \Pi(x, y) = x \quad (10)$$

$$B \overset{\mathbf{f}}{\circ} C = C, \text{ when } f_i(x, y) = \mathcal{D}_2 M(x, y) = \mathbb{I}_{\{y \leq x\}} \quad (11)$$

$$B \overset{\mathbf{f}}{\circ} C = \bar{C}, \text{ when } f_i(x, y) = \mathcal{D}_2 W(x, y) = \mathbb{I}_{\{1-x \leq y\}} \quad (12)$$

where \bar{C} denotes the survival copula of C .

More precisely, let (V_1, \dots, V_n) and (U_1, \dots, U_n) be two independent random vectors with joint distribution functions B and C , respectively. The survival copulas of B and C are denoted by \bar{B} and \bar{C} , and are defined by

$$\bar{B}(x_1, \dots, x_n) = \mathbb{P}(1 - V_i \leq x_i, i = 1, \dots, n),$$

and

$$\bar{C}(x_1, \dots, x_n) = \mathbb{P}(1 - U_i \leq x_i, i = 1, \dots, n).$$

Property 3 (Linearity). The composition operation is linear in both component copulas:

$$(\lambda B_1 + (1 - \lambda) B_2) \overset{\mathbf{f}}{\circ} C = \lambda (B_1 \overset{\mathbf{f}}{\circ} C) + (1 - \lambda) (B_2 \overset{\mathbf{f}}{\circ} C) \quad (13)$$

$$B \overset{\mathbf{f}}{\circ} (\lambda C_1 + (1 - \lambda) C_2) = \lambda (B \overset{\mathbf{f}}{\circ} C_1) + (1 - \lambda) (B \overset{\mathbf{f}}{\circ} C_2) \quad (14)$$

for $\lambda \in [0, 1]$.

2.5 Convergence Properties

The composite copula framework exhibits convergence behaviour, as shown by [1]:

Theorem 3 (Uniform Convergence of Composite Copulas). Suppose that $\mathbf{f} = (f_1, \dots, f_n) \in \mathcal{F}$.

1. If the sequence of copulas $\{B_k\}_{k=1}^{\infty}$ converges to a copula B uniformly as $k \rightarrow \infty$, then the sequence of multivariate composite copulas $\{B_k \overset{\mathbf{f}}{\circ} C\}_{k=1}^{\infty}$ converges to $B \overset{\mathbf{f}}{\circ} C$ uniformly.
2. If the sequence of copulas $\{C_k\}_{k=1}^{\infty}$ converges to a copula C uniformly as $k \rightarrow \infty$, then the sequence of multivariate composite copulas $\{B \overset{\mathbf{f}}{\circ} C_k\}_{k=1}^{\infty}$ converges to $B \overset{\mathbf{f}}{\circ} C$ uniformly.

Theorem 4 (Convergence with Respect to Function Vectors). Let

$\{\mathbf{f}_k(x, y) = (f_{1,k}(x, y), \dots, f_{n,k}(x, y))\}_{k \geq 1}$ be a sequence of function vectors with $\mathbf{f}_k \in \mathcal{F}$ for each $k \geq 1$, and let $\mathbf{f} = (f_1, \dots, f_n) \in \mathcal{F}$.

If for each $i = 1, \dots, n$, the sequence of functions

$$\left\{ \int_0^1 |f_{i,k}(x, y) - f_i(x, y)| dy, x \in [0, 1] \right\}_{k \geq 1} \quad (15)$$

converges uniformly to zero as $k \rightarrow \infty$, then the sequence of multivariate composite copulas $\{B \overset{\mathbf{f}_k}{\circ} C\}_{k \geq 1}$ converges uniformly to the multivariate composite copula $B \overset{\mathbf{f}}{\circ} C$.

3 NUMERICAL ILLUSTRATION OF COMPOSITE COPULAS

This section provides numerical illustrations of the key characteristics of multivariate composite copulas, including convergence behaviour, reproduction properties, and continuity. These demonstrations serve to validate the theoretical properties established in Section 2 and provide intuition for the practical behaviour of composite copulas.

3.1 Convergence Behaviour of Composite Copulas

The convergence properties of composite copulas can be illustrated through the Frank copula family, which exhibits interesting limiting behaviour. Recall that the bivariate Frank copula is given by [6]:

$$C(u, v; \gamma) = -\frac{1}{\gamma} \ln \left(1 + \frac{(e^{-\gamma u} - 1)(e^{-\gamma v} - 1)}{e^{-\gamma} - 1} \right) \quad (16)$$

$$\gamma \in \mathbb{R} \setminus \{0\}$$

γ	Kendall's τ	Spearman's ρ	γ	Kendall's τ	Spearman's ρ
-400	0.3238	0.4670	400	0.3237	0.4670
-100	0.3244	0.4679	100	0.3243	0.4677
-10	0.3883	0.5498	10	0.3884	0.5501
-0.01	0.6631	0.8455	0.01	0.6631	0.8455

Table 1: Kendall's τ and Spearman's ρ for Composite Copula $B \overset{f}{\circ} C$ with Varying γ

The Frank copula has the property that:

- Fréchet-Hoeffding upper bound

$$\lim_{\gamma \rightarrow \infty} C(u, v; \gamma) = M(u, v) \tag{17}$$

- Fréchet-Hoeffding lower bound

$$\lim_{\gamma \rightarrow -\infty} C(u, v; \gamma) = W(u, v) \tag{18}$$

- Independence/Product copula

$$\lim_{\gamma \rightarrow 0} C(u, v; \gamma) = \Pi(u, v) \tag{19}$$

3.1.1 Convergence Demonstration

We demonstrate the convergence behaviour using the following setup:

- Component copula B : Gumbel copula with parameter $\alpha = 3$
- Component copula C : Clayton copula with parameter $\beta = 1$
- Function vector: $f_1(x, y) = f_2(x, y) = \mathcal{D}_2 D(x, y; \gamma)$, where D is a Frank copula with parameter γ

According to Property 2.4 and Theorem 3, we expect:

$$\lim_{\gamma \rightarrow \infty} B \overset{f}{\circ} C = C \tag{20}$$

$$\lim_{\gamma \rightarrow -\infty} B \overset{f}{\circ} C = \bar{C} \tag{21}$$

$$\lim_{\gamma \rightarrow 0} B \overset{f}{\circ} C = B \tag{22}$$

3.2 Comprehensive Numerical Illustration

Measures of dependence. To help interpret the values reported in Table 1, we briefly recall two common rank-based measures of dependence.

Kendall's tau. For a bivariate copula C , Kendall's τ is defined by

$$\tau(C) = 4 \int_{[0,1]^2} C(u, v) dC(u, v) - 1.$$

Equivalently, if (X_1, Y_1) and (X_2, Y_2) are independent copies, then

$$\tau = \mathbb{P}((X_1 - X_2)(Y_1 - Y_2) > 0) - \mathbb{P}((X_1 - X_2)(Y_1 - Y_2) < 0).$$

Intuitively, Kendall's τ measures the difference between the probability of concordance and discordance, taking values in $[-1, 1]$.

Spearman's rho. For a copula C , Spearman's ρ is defined by

$$\rho(C) = 12 \int_0^1 \int_0^1 (C(u, v) - uv) dudv.$$

It can also be interpreted as the correlation between the ranks of the variables. Like Kendall's τ , it takes values in $[-1, 1]$, where 1 corresponds to perfect positive dependence, -1 to perfect negative dependence, and 0 to independence.

3.3 Interpretation of Numerical Results

Following [1], a reproduction of the numerical illustrations are shown above. We use the fact that the Frank copula is stochastically ordered that its limits coincide with the Fréchet-Hoeffding upper and lower bounds. Let bivariate copula B be Gumbel copula with parameter $\alpha = 3$, C be Clayton copula with parameter $\beta = 1$ and D be Frank copula with parameter γ . Then choose functions $f_1(x, y) = f_2(x, y) = \mathcal{D}_2 D(x, y; \gamma)$. This configuration creates a composite copula $B \overset{f}{\circ} C$ whose dependence structure varies continuously with γ .

The component copulas exhibit distinct dependence characteristics: the Gumbel copula B demonstrates strong upper dependence with Kendall's $\tau = 0.6667$ and Spearman's $\rho = 0.8482$, while the Clayton copula C shows strong lower tail dependence with $\tau = 0.3333$ and $\rho = 0.4785$. The survival Clayton copula \bar{C} shares the same dependence measures as C but with reversed tail dependence patterns.




Figure 1: Convergence of Composite Copula $B^f \circ C$ with varying γ parameter.

The observed patterns in (a) - (c) reveal the underlying mechanism: as $\gamma \rightarrow \infty$, the Frank copula approaches the Fréchet-Hoeffding upper bound, causing the composite construction to emphasize the characteristics of Clayton copula C , as we expected from equation 22. Kendall's τ and Spearman's ρ also show a decreasing pattern, in which it shows a weaker concordance and converging toward the Clayton copula values. As we can see, the upper tail dependence gradually disappears as γ increases, remains only the lower tail dependence of Clayton copula.

Similarly, as γ approaches 0, the Frank copula corresponds to the Independent copula, an emphasis of Gumbel copula B is expected from equation 22, presented in (g) - (i). The dependence measures increase to $\tau \approx 0.663$ and $\rho \approx 0.845$, closely matching the theoretical Gumbel values, while the scatter plots show the development of strong upper tail dependence charac-

teristic of Gumbel copulas.

The survival copula \bar{C} follows the same mechanism as copula C , but in this case the Frank copula corresponds to Fréchet-Hoeffding lower bound when γ approaches to $-\infty$. And instead of upper tail dependence, its lower tail dependence disappears as γ is getting smaller.

4 EMPIRICAL ANALYSIS OF GERMAN GOVERNMENT BOND YIELDS

This section presents a quick overview of the reproduction conducted using the Chinese treasury bond dataset used in the paper [1]. And an original empirical analysis of the dependence structure between short-term and long-term German government bond yields. The dataset and preliminary dependency analysis provide the foundation for our copula modeling approach.

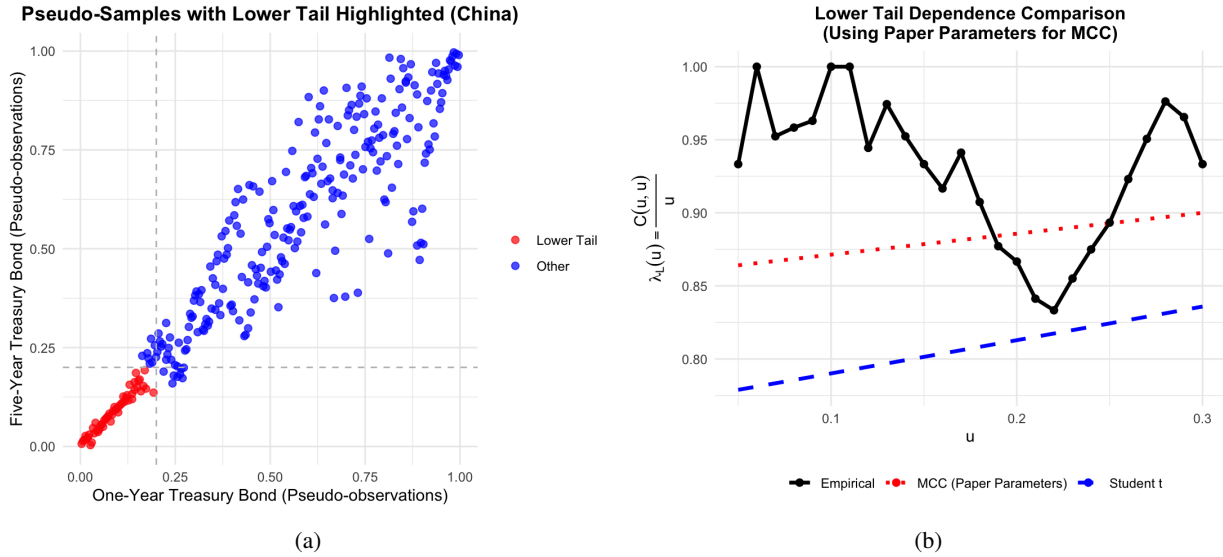


Figure 2: (a) Pseudo-Samples with Lower 20% Tail Highlighted (Chinese Treasury Bonds). (b) Reproduction of the quantile lower dependence comparison from the original paper [1] Where MCC stands for multivariate composite copula with the estimated parameter vector $\hat{\theta} = (\hat{\xi}_1, \hat{\xi}_2, \hat{\gamma}) = (0.97513, 0.40869, -0.01474)$, where $\hat{\xi}_1$ and $\hat{\xi}_2$ are the parameters for two student-t copulas as B and C respectively, $\hat{\gamma}$ is the parameter for the Gaussian copulas D_1 and D_2 .

4.1 Chinese Treasury Bond Analysis

We begin by reproducing the key empirical patterns from the original study, which examined Chinese treasury bonds. Figure 2a shows the pseudo-observations with highlighted tail regions, while Figure 2b displays the tail dependence measures for Chinese government bonds. Data can be found in [7].

These reproduced figures from the original study provide a benchmark for our analysis of German government bonds and demonstrate the methodological consistency across different bond markets.

4.2 German Data Description

We analyze monthly German government bond yields from January 1960 onwards from [4], comprising 790 observations. The dataset includes:

- **3-month rate:** Short-term government bond yield
- **10-year rate:** Long-term government bond yield

The first few observations of our dataset are presented below:

Date	3-Month Rate	10-Year Rate
1960-01-01	4.36	6.50
1960-02-01	4.47	6.40
1960-03-01	4.71	6.40
1960-04-01	4.59	6.40
1960-05-01	4.64	6.40
1960-06-01	5.25	6.60

Table 2: First Six Observations of German Government Bond Yields.

Summary statistics reveal the following characteristics:

Table 3: Summary Statistics of German Government Bond Yields

Statistic	3-Month Rate	10-Year Rate
Mean	4.27	5.30
Standard Deviation	3.14	2.86

The Kendall’s τ correlation between the 3-month and 10-year rates is 0.757, indicating strong positive dependence between short-term and long-term yields.

4.3 Pseudo-Samples Transformation

To analyze the dependence structure independent of marginal distributions, we transform the observed yields to pseudo-observations using the empirical distribution function:

$$U_i = \frac{\text{Rank}(X_i)}{n+1}, \quad V_i = \frac{\text{Rank}(Y_i)}{n+1} \quad (23)$$

where X_i and Y_i represent the 3-month and 10-year yields, respectively, and $n = 790$ is the sample size, $i = 0, 1, \dots, n$. The resulting pseudo-samples range from 0.0013 to 0.9988 for both series.

The scatter plot of pseudo-observations (Figure 3) reveals the dependence structure between the transformed variables.

Pseudo-Samples with Lower Tail Highlighted (Germany)




Figure 3: Pseudo-observations with Lower 20% Tail Highlighted (German Government Bonds).

4.3.1 Empirical Lower Tail Dependence at 20% Quantile

To quantify the lower tail dependence at a specific threshold, we analyze the joint behaviour in the lower 20% quantile. Our analysis reveals:

- **Points in lower 20% quantile:** 135 joint observations
- **Lower tail ratio:** 0.8544

This lower tail ratio is calculated as:

$$\lambda_L^{emp}(0.2) = \frac{N}{n \times 0.2} \quad (24)$$

where $n = 790$ is the total number of observations, and N is number of points where both $U \leq 0.2$ and $V \leq 0.2$. A ratio of 0.8544 indicates strong positive lower tail dependence, meaning that when one yield falls into its lowest 20%, the other yield has an 85.44% probability of also being in its lowest 20%. This substantially exceeds what would be expected under independence (which would yield a ratio of 0.2) and demonstrates the tendency for German short-term and long-term yields to experience joint downturns.

4.4 Tail Dependence Analysis

4.4.1 The empirical quantile dependence

We now introduce the formal definition of the lower tail ratio. We call it the empirical lower tail dependence function $\lambda_L^{emp}(u)$ measures the probability that both series experience extremely low values simultaneously:

$$\lambda_L^{emp}(u) = P(U \leq u | V \leq u) = \frac{P(U \leq u, V \leq u)}{P(V \leq u)} \quad (25)$$

Similarly, the upper tail dependence function $\lambda_U^{emp}(u)$ is defined as:

$$\lambda_U^{emp}(u) = P(U > u | V > u) = \frac{P(U > u, V > u)}{P(V > u)} \quad (26)$$

Our analysis of the lower 20% quantile reveals 135 observations, with a lower tail dependence of 0.854, indicating strong positive lower tail dependence.

4.4.2 Empirical Tail Dependence Functions

Table 4 presents the empirical tail dependence functions across different quantile thresholds:

Threshold(u)	$\lambda_L^{emp}(u)$	$\lambda_U^{emp}(u)$
0.01	0.1429	—
0.10	0.9241	—
0.12	0.9787	—
0.20	0.8544	—
0.80	—	0.7289
0.90	—	0.5952
0.96	—	0.7419
0.99	—	0.0000

Table 4: Detailed Empirical Tail Dependence

Key findings from our tail dependence analysis:

- **Average Lower Tail Dependence:** 0.7854 (strong)
- **Average Upper Tail Dependence:** 0.6202 (moderate)
- **Maximum Lower Tail Dependence:** 0.9787 (near-perfect dependence)
- **Maximum Upper Tail Dependence:** 0.7419

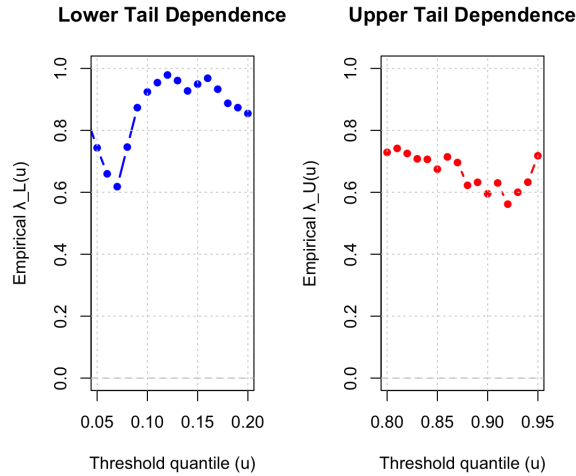


Figure 4: Pseudo-observations with Lower Tail Highlighted (German Government Bonds)

4.5 Benchmark Comparison Analysis

The comparison with independence benchmarks gives an insight on the economic significance of the tail dependence:

Scenario	Observed	Independence	Ratio
Lower tail ($u=0.20$)	0.8544	0.2000	4.27
Upper tail ($u=0.80$)	0.7289	0.2000	3.64
Lower tail ($u=0.01$)	0.1429	0.0100	14.29
Upper tail ($u=0.99$)	0.0000	0.0100	0.00

Table 5: Comparison with Independence Benchmark

The independence benchmarks serve as a natural reference point because they represent the minimum meaningful threshold for dependence analysis. They also provide scale-invariant measures of dependence strength.

5 LIMITATIONS OF STANDARD COPULA FAMILIES

This section explores the insufficiencies of traditional copula families for capturing the asymmetric tail dependence structure observed in German government bond yields. We fit four common copula models and demonstrate their limitations through parameter estimation, tail dependence analysis, and formal goodness-of-fit testing.

5.1 Baseline Copula Specifications

We consider four standard copula families often used in finance, representing different dependence struc-

tures:

- **Gaussian Copula:** Widely used to model the dependence structure in finance and insurance.
- **Student-t Copula:** Captures symmetric tail dependence with equal lower and upper tail dependence, again usually used in finance modelling, especially risk management.
- **Clayton Copula:** Captures lower tail dependence only, with no upper tail dependence. Often used in finance and environmental sciences.
- **Gumbel Copula:** Captures upper tail dependence only, with no lower tail dependence. Used in finance, risk management and time series analysis.

All models are fitted to the pseudo-observations (U_i, V_i) using the maximum pseudo-likelihood method.

Maximum pseudo-likelihood estimation. The copula models are fitted to the pseudo-observations (U_i, V_i) using the *maximum pseudo-likelihood* (MPL) method. Here, $(U_i, V_i) \in [0, 1]^2$ are obtained by transforming the original data using their empirical marginal distribution functions, so that the marginals are approximately uniform on $[0, 1]$.

Let $c_\theta(u, v)$ denote the copula density corresponding to a copula C_θ with parameter vector θ . Given a sample $\{(U_i, V_i)\}_{i=1}^n$, the (pseudo-)log-likelihood is defined as

$$\ell(\theta) = \sum_{i=1}^n \log c_\theta(U_i, V_i).$$

The maximum pseudo-likelihood estimator is then

$$\hat{\theta} = \arg \max_{\theta} \ell(\theta).$$

Interpretation. Each term $c_\theta(U_i, V_i)$ represents how likely the model with parameter θ considers the observation (U_i, V_i) . The log-likelihood $\ell(\theta)$ aggregates this information over all observations, and the estimator $\hat{\theta}$ selects the parameter values that make the observed data as likely as possible under the chosen copula model.

The term “pseudo” reflects the fact that the true marginal distributions are unknown and are replaced by their empirical counterparts. As a result, the likelihood is not a full likelihood for the joint distribution, but rather a likelihood based only on the dependence structure captured by the copula.

Connection to the implementation. In the code, this procedure is carried out using the function `fitCopula(..., method = "ml")` or `"mpl"`. For Archimedean copulas (such as Clayton and Gumbel),

the MPL method is typically used, while for elliptical copulas (Gaussian and t), maximum likelihood is applied directly to the pseudo-observations. In both cases, the optimization routine numerically maximizes the log-likelihood $\ell(\theta)$.

5.2 Parameter Estimation Results

Table 6 presents the maximum likelihood estimates for each copula specification:

Copula	Parameter 1	Degrees of Freedom	AIC	BIC
Gaussian	$\rho = 0.9131$	—	-1414.175	-1409.503
Student-t	$\rho = 0.9152$	$\nu = 11.67$	-1418.060	-1408.716
Clayton	$\alpha = 6.2321$	—	-1136.681	-1132.009
Gumbel	$\alpha = 3.2964$	—	-1242.573	-1237.901

Table 6: Parameter Estimates for Baseline Copula Models

The Gaussian and Student-t copulas achieve the lowest AIC and BIC values, suggesting better overall fit to the dependence structure. However, as we will demonstrate, this masks critical deficiencies in capturing the observed tail behavior.

5.3 Tail Dependence Analysis

The fundamental limitation of standard copula families becomes apparent when examining their implied tail dependence, presented in Table 7:

Table 7: Tail Dependence of Baseline Copulas

Copula	$\lambda_L^{emp}(0.2)$	$\lambda_U^{emp}(0.8)$
Gaussian	0.7726	0.7696
Student-t	0.7876	0.7740
Clayton	0.8916	0.6683
Gumbel	0.6842	0.7948
Empirical	0.8544	0.7289

Each copula family fails to capture the empirical tail dependence structure in distinct ways:

- **Gaussian Copula:** Shows nearly symmetric tail dependence ($\lambda_L^{emp} = 0.7726$ and $\lambda_U^{emp} = 0.7696$).
- **Student-t Copula:** Shows nearly symmetric tail dependence ($\lambda_L^{emp} = 0.7876$ and $\lambda_U^{emp} = 0.7740$), contradicting the empirical asymmetry where lower tail dependence is stronger than upper tail dependence.
- **Clayton Copula:** Captures strong lower tail dependence of 0.8916, and an upper tail dependence of 0.6683. Presents the asymmetry but not close enough to the empirical ones.
- **Gumbel Copula:** Captures strong upper tail dependence of 0.7948, and a moderate lower tail dependence of 0.6842. Does not match the pattern of empirical data.

No single standard copula can simultaneously capture both the strong lower tail dependence (0.8544) and moderate upper tail dependence (0.7289) observed empirically.

5.4 Visual Examination of Simulated Data

Figure 5 displays simulated data from each fitted copula alongside the empirical pseudo-observations:

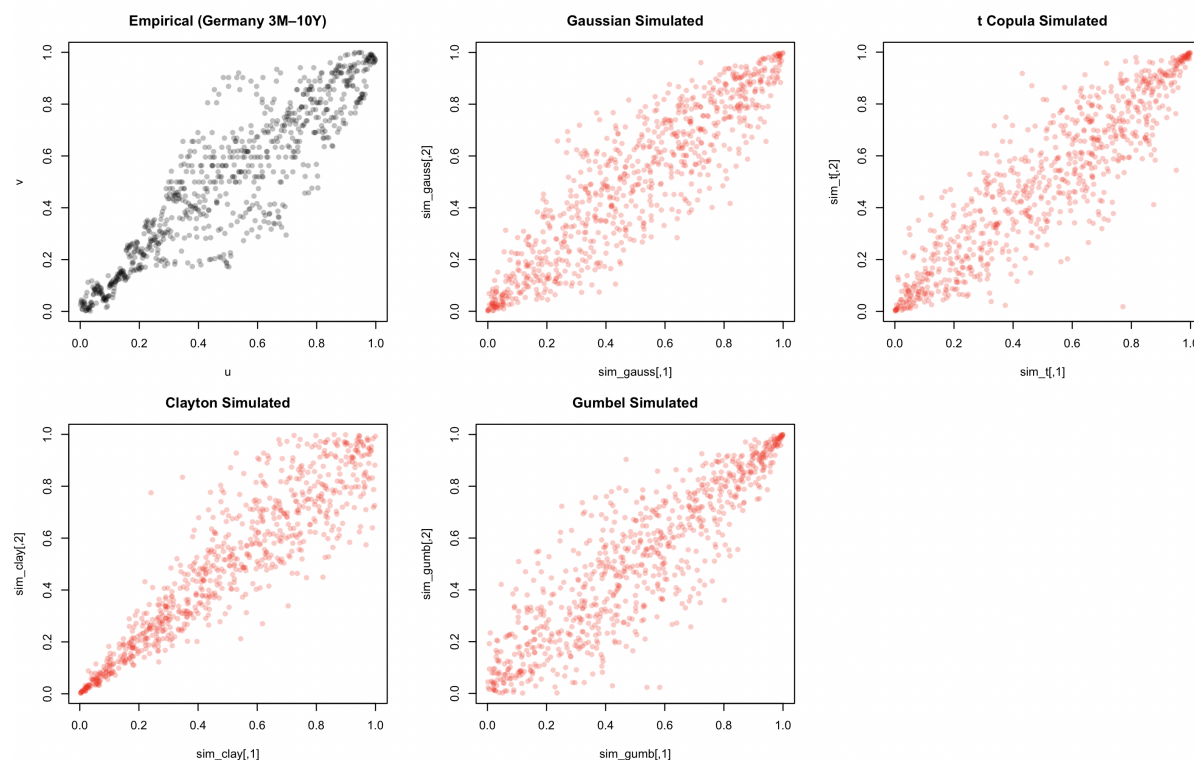


Figure 5: Comparison of Empirical Data and Copula Simulations

Visual inspection reveals that:

- **Gaussian simulations** fail to capture the concentration of points in the lower-left quadrant
- **Student-t simulations** show symmetric tail clustering that doesn't match the empirical asymmetry
- **Clayton simulations** correctly capture lower tail clustering but miss upper tail dependence
- **Gumbel simulations** correctly capture upper tail clustering but miss lower tail dependence

5.5 Theoretical Implications

The systematic failure of standard copula families reveals important theoretical limitations:

- **Symmetry Constraint:** Symmetry Constraint: Elliptical copulas (Gaussian copulas, t copulas) enforce symmetric dependencies, as mentioned by [6]. Further more, they are empirically invalid in yield curve dynamics, since it is unlikely that the German bond yields exhibit symmetric dependencies.
- **Single-Tail Specialization:** One-Tail Specialization: Some Archimedean copulas, like the

ones we are testing (Clayton copulas, Gumbel copulas). They specialize only in one tail (upper or lower) and fail to capture the two-tailed dependencies often observed in financial markets.

These findings provide compelling justification for exploring composite copula models that can:

1. Separately parameterize lower and upper tail dependence
2. Capture asymmetry in tail behaviour
3. Maintain computational tractability
4. Provide economically interpretable parameters

6 IMPLEMENTATION OF THE COMPOSITE COPULA APPROACH

Building on the theoretical and empirical analyses in the preceding sections, this section presents the implementation of the composite copula method for modelling the dependence structure of German government bond yields.

6.1 Model Construction and Estimation Strategy

Following the framework of [1], a bivariate composite copula is constructed as

$$B \circ^f C(u_1, u_2) = \mathbb{E}[B(f_1(u_1, U_1), f_2(u_2, U_2))],$$

where $(U_1, U_2) \sim C$ and the function vector $f = (f_1, f_2)$ is induced by a third copula D .

This construction allows the model to combine different dependence mechanisms: B primarily controls one aspect of dependence (e.g., lower tail), C introduces additional dependence structure, and D governs the interaction between them.

Estimation strategy. Although [1] propose a likelihood-based estimation procedure, it is computationally demanding due to the nested structure of the model.

Instead, we adopt an iterative simulation-based approach:

1. Specify candidate copula families for (B, C, D) ,
2. Select candidate parameter values over a reasonable range,
3. Simulate samples from the resulting composite copula,
4. Compare simulated samples with the empirical pseudo-observations using:
 - scatter plots (global dependence structure),
 - contour plots (density concentration),
 - tail dependence measures.

The final model is selected as the configuration that best reproduces both the visual structure and the empirical dependence statistics.

6.3 Final Model and Empirical Fit

Based on the above exploration, the following configuration was selected:

$$B = \text{Clayton}(10), C = t(0.85, \nu = 2), D = \text{Normal}(0.52).$$

Interpretation. This combination reflects the empirical dependence structure observed in Section 5:

- The Clayton copula ($\theta = 10$) captures the strong lower tail and moderate upper tail dependence.
- The t -copula introduces heavy-tailed symmetric dependence, allowing flexibility in joint extremes.

6.2 Parameter Search and Model Selection

To identify an appropriate composite copula specification, we explored a range of candidate models by varying both copula families and parameter values.

Search strategy. The selection process focused on three key aspects:

- **Lower tail dependence:** matching the strong clustering observed for $u \leq 0.2$,
- **Upper tail dependence:** reproducing the moderate dependence for $u \geq 0.8$,
- **Global structure:** maintaining realistic dependence in the central region.

Rejected configurations. Several candidate models were tested but rejected:

- **Weaker Clayton parameters** (e.g., $\theta = 5$) underestimated lower tail clustering.
- **Higher degrees of freedom in the t -copula** reduced tail dependence, producing overly Gaussian-like behaviour.
- **Lower Gaussian correlation** in the linking copula failed to reproduce the central dependence structure.

In Figure 6, The rejected models illustrate how different parameter choices fail to capture key features of the data. A weaker Clayton parameter underestimates upper tail clustering, a higher t -copula degree of freedom reduces tail dependence, and a lower Gaussian correlation tighten the central dependence structure.

- The Gaussian copula provides smooth dependence in the central region and stabilizes the overall structure.

Fit to empirical data. Simulations from the composite copula show strong agreement with the empirical pseudo-observations:

- Lower tail dependence: $\lambda_L^{\text{composite}} \approx 0.8544$, matching the empirical value,
- Upper tail dependence: $\lambda_U^{\text{composite}} \approx 0.7089$, close to the empirical value 0.7289.

In contrast to standard copulas, the composite model successfully captures both the asymmetry and magnitude of tail dependence.

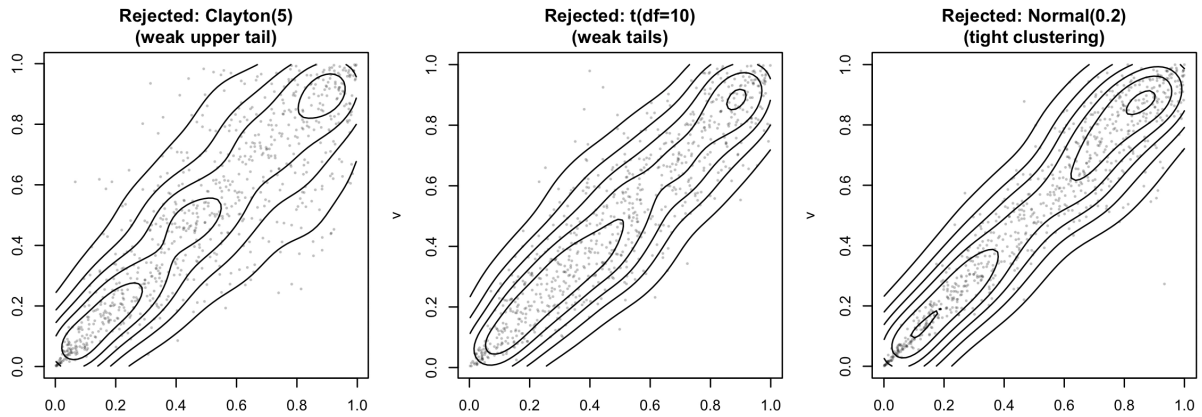


Figure 6: Comparison of candidate composite copula models. The rejected models fail to capture one or more key features.

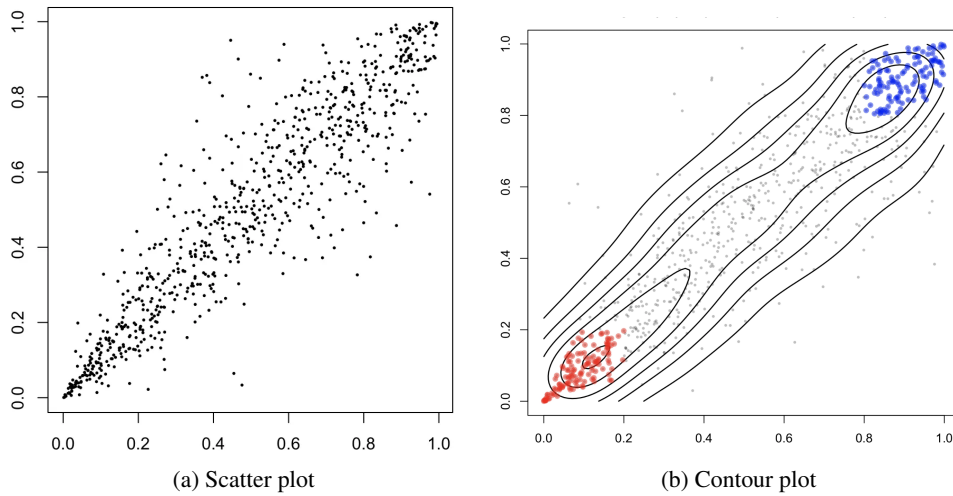


Figure 7: Simulated distribution of the fitted composite copula using $B = \text{Clayton}(10)$, $C = t(0.85, 2)$, and $D = \text{Normal}(0.52)$.

Figure 7 confirms that the simulated distribution reproduces the key features of the empirical data, including strong lower tail clustering and moderate upper tail dependence.

section successfully captures the key dependence features observed in German bond yield data, offering a flexible but not too complex alternative to standard parametric copulas.

6.4 Implementation Considerations

Overall, the computational cost of composite copulas in the binary case is manageable. Further study can be conducted to extend it to higher dimensions, which may require numerical integration, dimensionality reduction techniques, or parallel simulation schemes. Approximate likelihood methods (e.g., surrogate likelihood methods, variational methods) can be explored to improve parameter estimation efficiency.

Summary. The fitted composite copula shown in this

7 CONCLUSION

The report reviews the construction of the multivariate composite copula. We begin by restating the mathematical foundations of the composite copulas, and we also test several properties, like the reproduction property, by generating numerical illustrations in Section 3. We apply them to the dependence structure of German government bond yields. During this part, we first examine the empirical pseudo-sample of the data, capture its strong lower tail dependence and moderate upper tail dependence. Based on this, we try to fit four

different standard copulas using goodness-of-fit diagnostics, but none of them captures the full dependence structure.

To address this issue, we implemented the composite copula that follows the proposed construction by the [1], so that the tail dependence is controlled by three copulas, Clayton, Student-t and Normal. With each parameter carefully chosen, we successfully produced a well-fitting model that captures both tail dependence in a more accurate strength than the standard four copulas.


Overall, the results confirm the theoretical foundation of the multivariate composite copula; it also justifies the usefulness in certain fields like economics, where the composite copula may perform better than the standard copulas. Future work may extend these ideas to higher-dimensional settings, develop faster parameter estimation techniques, or explore applications in financial risk management and multivariate time series modelling.

REFERENCES

- [1] Xie, J., Fang, J., Yang, J. & Bu, L. Multivariate Composite Copulas. *ASTIN Bulletin*. **52**, 145-184 (2021)
- [2] Sklar, M. Fonctions de répartition à n dimensions et leurs marges. *Publications De L'Institut De Statistique De L'Université De Paris*. **8** pp. 229-231 (1959)
- [3] Nelsen, R. An Introduction to Copulas. (Springer Science & Business Media,2006)
- [4] Organization for Economic Co-operation and Development 3-Month or 90-day Rates and Yields: Interbank Rates for Germany. (Federal Reserve Economic Data, Federal Reserve Bank of St. Louis,2024), <https://fred.stlouisfed.org/series/IR3TIB01DEM156N>
- [5] Organization for Economic Co-operation and Development Long-Term Government Bond Yields: 10-year: Main (Including Benchmark) for Germany. (Federal Reserve Economic Data, Federal Reserve Bank of St. Louis,2024), <https://fred.stlouisfed.org/series/IRLTLT01DEM156N>
- [6] Embrechts, P., Lindskog, F. & McNeil, A. Modelling Dependence with Copulas and Applications to Risk Management. (Department of Mathematics, ETH Zürich,2001,9), <https://www.math.ethz.ch/finance>, Working paper
- [7] ChinaBond ChinaBond Yield Curves. (2025), https://yield.chinabond.com.cn/cbweb-main/yield_main?locale=en_US, Accessed: 2025-11-20

INTERVIEW WITH ROBERT ROBERE

Aahaan Rawal and Aditya Sharma



$\delta\epsilon$: Could you briefly introduce yourself, your academic position, and where you're from?

My name is Robert Robere. I'm an assistant professor here in the Computer Science department in the School of Computer Science. I'm physically originally from Newfoundland. Academically, I did my undergrad in Newfoundland. I did my master's and PhD at University of Toronto, and then I did post-docs at UC Berkeley, Rutgers and Institute for Advanced Study. Then I came here to McGill.

$\delta\epsilon$: Could you describe your research interest in a general sense, something that's accessible to an undergrad, and a specific question that you're interested in right now?

Broadly speaking, what I'm interested in is computational complexity theory. Particularly what I'm interested in is improving lower bounds. People have maybe heard of this P vs. NP problem, and while I'm not exactly focused on proving P vs. NP, I'm more focused on proving lower bounds. These are trying to understand the sorts of techniques that are needed to attack big separation questions in complexity theory. In particular, an area that I've worked a lot in is called proof complexity. This has a lot of links with areas like

cryptology and logic, and it studies like the lengths of proofs in particular.

But in general, I kind of like thinking about problems that have both a complexity or an algorithm/proof type feeling to them, and also I like problems with some interesting combinatorial content. A lot of lower bound problems boil down to proving an interesting combinatorial statement, and I like thinking about those kinds of problems too.

$\delta\epsilon$: Do you have any advice for undergraduate students studying CS?

I would say that we are at a very disruptive point with computer science right now, particularly with the development of LLMs, and we literally have no idea both technologically, scientifically, politically, or economically where it's going to go. Maybe this is the singularity, who knows, right?

That being said, I still think that there is a huge opportunity for people to study computer science and mathematics. Because the human element is never, ever going to get removed from it. In my own experience, the most valuable way to use an AI tool is to use it to give feedback on my own stuff. Don't ask it to answer questions, ask it to give you feedback on what you give it. That's a very valuable tool, using it as a sounding board. It's very, very good at that, and it's very, very good at literature search. It's very disruptive in a lot of ways, but it is not going to replace human elements.

So to an undergrad, I would honestly say: learn how to use the tools, but you've got to include yourself in the loop. Because that's really where the value comes from; the value doesn't come from the tool by itself. Everyone who says it's going to replace people is wrong, because there's always going to be a human in a loop; humans are the basis of everything that we do. Understanding the fundamentals, the systems, how they are developed, learning how to use them in a way that makes you more effective, these are all valuable things to do. Don't worry about the hype. Don't give away your soul, I guess, would also be my big advice.

$\delta\epsilon$: Which probably would like to see the most to be solved in the next decade?

Here's a question that's related to neural nets that I think is super interesting. So, a neural net is like a layered circuit where every single gate in it is like a threshold gate where you check the weights and then you do some sort of softmax thing and then you out-

put whether it activates or not. If I take a neural net of depth one, that's just a threshold gate. It's very easy to prove that some things can't be represented that way. If I take a neural net of depth two, we can prove quadratic lower bounds on the size. So neural nets of depth two can represent every function. But, some of these representations, the best thing we know of is exponentially large, just by doing some sort of truth table, right? And the best lower bound that we have is n^2 . So we only know how to prove a lower bound for a function on n bits with n^2 neurons in it.

So I think that it is very likely that major progress is going to be made on this in the next 10 years. And it would be absolutely amazing.

$\delta\epsilon$: You mentioned that your PhD advisor was Stephen Cook, who is very famous, I would say, for undergrads. Could you describe your experience kind of being mentored and advised by him?

Steve rules. I met Steve in a very roundabout way. So when I was an undergrad, my undergraduate honours advisor, Antonina, was also Steve's PhD student, and she convinced him to come visit, to give a plenary lecture. I was doing USRA type stuff with her as an undergrad, and I was working on stuff that Steve was working on, because she had said that "oh, like, I was talking to Steve and they're working on this problem and it seems approachable. He said, if there's any students here who'd be interested, you should work on it." So I started working on it, and then I proved some things about it.

When he visited, I ended up spending the whole visit talking to him. At the end of the visit, he said "oh, you should come to Toronto". And I was like, "okay, but I'm not supposed to graduate for like 2 years." And he was like, "oh, well, it's probably better if you came sooner"; I think because he was getting older and thinking about retiring. So I did a complicated thing where I dropped a major to a minor and graduated early, because if Steve tells you to go to Toronto, you go.

I worked with him really closely during my masters. He was really my primary advisor during my masters. And then in my PhD I started working more on complexity lower bound type problems, so I started talking to Toni more. And Steve was getting towards retirement age. So, over the course of my PhD, I moved from working with him primarily to working with Toni primarily. But he was very, very active. We published a paper together while I was a student. He's just like a lovely person. He's an angel of a human being.

So it was a really, really good experience, actually. I was very lucky in my choice of advisors. I think if I can give one piece of advice to anyone who wants to go to grad school is get a good advisor. I got extremely lucky with mine, they were both outstanding.

$\delta\epsilon$: Do you think breakthroughs are more likely to come from new techniques or rethinking old ones?

I think it's actually pretty rare that there is some unsolvable problem that people have tried forever to solve and they're like: "I don't know, don't have it". And then some person comes out of nowhere and is like: "here's my brand new theory that just resolves everything". This almost never, ever, ever happens. It's always building upon something. And like I said earlier, usually what happens when a problem is sort of ripe to be solved, is that the field around the problem has built up enough tools, reductions, and simplifications that they've broken this problem down to its nugget that you have to solve.

Once you get to that nugget, you need some new ideas. So that's where the new ideas usually really, really come in handy. But it's never in isolation, and that's why I said it's always modifications. Because you need that foundation of other reductions, known techniques, and counter examples; examples where this technique should work, but just because it barely doesn't work, that might give me some insight as to what the right way to proceed is. Honestly, it's a mix. It's almost never someone out of the blue just coming up with a brand new technique that just knocks it out of the park.

PROOF OF RIEMANN MAPPING THEOREM

Nitya Khirwar

This paper presents a comprehensive proof of the Riemann Mapping Theorem, a result in complex analysis first stated by Bernhard Riemann in 1851. The theorem asserts that any connected proper subdomain of the complex plane is conformally equivalent to the open unit disk. We follow the classical function-theoretic approach, systematically developing the necessary tools and concepts. Beginning with the Schwarz Lemma and the automorphism group of the unit disk, we establish the hyperbolic metric and its contraction properties under holomorphic mappings. The theory of normal families, culminating in Montel's theorem, provides the compactness principle required to solve an extremal problem: maximising the derivative at a fixed point within a family of bounded univalent functions. The proof leverages the square root function's expanding behaviour under the hyperbolic metric, which, when applied to a candidate extremal map that is not surjective, yields a contradiction. Thus, the extremal function must be a conformal bijection onto the unit disk, completing the construction of the Riemann map. The exposition is self-contained, assuming only standard undergraduate-level complex analysis, and includes detailed illustrations of the relevant geometric concepts.

1 INTRODUCTION

In complex analysis, two domains in the field of complex numbers \mathbb{C} are considered conformally equivalent if there is a bijective holomorphism between them with a holomorphic inverse. In 1851, Bernhard Riemann stated in his PhD thesis that any simply connected domain Ω in \mathbb{C} that is not all of \mathbb{C} is conformally equivalent to the open unit disk $\mathbb{D} := \{z \in \mathbb{C} : |z| < 1\}$, in a result known as the Riemann Mapping Theorem. Furthermore, he proposed that this Riemann map is unique. The goal of this paper is to cover the proof of this theorem, closely following the corresponding sections of Theodore W. Gamelin's *Complex Analysis* [1].

The proof is set up as follows. In Section 2, we introduce the Schwarz Lemma (Theorem 1), a foundational constraint on holomorphic self-maps of \mathbb{D} , and use it to characterize the automorphism group $\text{Aut}(\mathbb{D})$ as the family of fractional linear transformations (Theorem 4). This leads naturally to Pick's Lemma (Theorem 5) and the *hyperbolic metric* on \mathbb{D} , an $\text{Aut}(\mathbb{D})$ -invariant metric under which every holomorphic self-map of \mathbb{D} is a contraction (Theorem 7). We also introduce the *spherical metric* on the extended complex plane \mathbb{C}^* , which governs notions of normal convergence in later sections.

In Section 3, we develop the theory of *normal families*. The key results are Hurwitz's Theorem (Theorem 8), which tracks the zeros of a normally convergent sequence, and a preservation result (Theorem 9) showing that normal limits of univalent functions remain univalent. Section 4 applies this theory via Montel's Theorem (Theorem 10), a compactness criterion guaranteeing that uniformly bounded families of holomorphic functions on Ω have normally convergent subsequences. As a consequence, we establish the existence of *extremal functions* – holomorphic functions on Ω that maximize $|f'(z_0)|$ at a fixed point z_0 within a bounded family (Theorem 13).

In Section 5, we formulate the Riemann map as the solution to this extremal problem by constructing a nice contradiction. The key geometric ingredient is the expanding property of the square root function under the hyperbolic metric (Lemma 14): if $\Omega \subsetneq \mathbb{D}$ is simply connected, one can always construct a conformal map $\psi : \Omega \rightarrow \mathbb{D}$ with $|\psi'(0)| > 1$, strictly increasing the derivative. Combining this with Montel's theorem, we show that the extremal function ϕ must map Ω onto all of \mathbb{D} , completing the proof of the Riemann Mapping Theorem (Theorem 17).

2 AUTOMORPHISMS OF THE UNIT DISK

To achieve the preliminary results that help us construct the unique Riemann map and complete the proof of the Riemann Mapping Theorem in later sections, we first need to introduce the group of automorphisms of the unit disk \mathbb{D} , which will be an essential ingredient of our map. For this, we require formal definitions for a conformal map and a conformal self-map of a domain in \mathbb{C} .

Definition. A *conformal map* between domains Ω_1 and Ω_2 is a bijective holomorphism $f : \Omega_1 \rightarrow \Omega_2$ with a holomorphic inverse. If $\Omega_1 = \Omega_2 = \Omega$, we call f a *conformal self-map* of Ω .

Definition. A *conformal self-map* of the domain Ω is a bijective holomorphism $f : \Omega \rightarrow \Omega$ with a holomorphic inverse.

Remark. These two notions are related but play distinct roles throughout the proof. A conformal map between two different domains Ω_1 and Ω_2 is the notion of equivalence we ultimately want to establish: the Riemann map itself will be a conformal map from Ω onto \mathbb{D} . A conformal self-map, by contrast, is a symmetry of a single domain. In this section, we study the group

$\text{Aut}(\mathbb{D})$ of conformal self-maps of \mathbb{D} in order to understand the internal geometry of \mathbb{D} via the hyperbolic metric. This geometric understanding is then used in Section 5, where we compose the candidate Riemann map with conformal self-maps of \mathbb{D} , particularly in Lemma 16, to derive a contradiction from the assumption that the extremal function does not map onto all of \mathbb{D} .

Theorem 1. (Schwarz Lemma). *Let f be a holomorphic function on \mathbb{D} . Suppose that $|f(z)| \leq 1$ for all $z \in \mathbb{D}$, and $f(0) = 0$. Then*

$$|f(z)| \leq |z|$$

on \mathbb{D} . Moreover, if at some $0 \neq z_0 \in \mathbb{D}$, equality holds (that is, $|f(z_0)| = |z_0|$), then $f(z) = \lambda z$ for some constant λ with $|\lambda| = 1$.

Proof. First note that since $f(z)$ has a root of order at least one at 0, $f(z) = zg(z)$ for some holomorphic function g on \mathbb{D} . Since $g(z) = \frac{f(z)}{z}$ is holomorphic, we let $r < 1$ and apply the maximum principle on the disk $C_r := \{|z| \leq r\}$. In $\partial C_r = \{|z| = r\}$, $|g(z)| = \frac{|f(z)|}{r} \leq \frac{1}{r}$. Therefore, by the maximum principle $|g(z)| \leq \frac{1}{r}$ everywhere in C_r . Taking the limit as $r \rightarrow 1$, we get that $\frac{|f(z)|}{|z|} \leq 1$ on \mathbb{D} .

If $|f(z_0)| = |z_0|$ for $z_0 \neq 0$ then $|g(z_0)| = 1$ and by the maximum principle, g is a constant, say λ . Therefore, $f(z) = \lambda z$ with $|\lambda| = 1$. \square

Corollary 2. *Let f be holomorphic on \mathbb{D} with $|f(z)| \leq 1$ and $f(0) = 0$ (with the same conditions as above). Then*

$$|f'(0)| \leq 1$$

and $|f'(0)| = 1$ if and only if $f(z) = \lambda z$ for some constant λ with $|\lambda| = 1$.

Proof. By the Cauchy integral formula,

$$|f'(0)| \leq \frac{1}{2\pi} \oint_{\mathbb{D}} \left| \frac{f(z)}{z^2} \right| dz \leq 1$$

Once again, consider $f(z) = zg(z)$ for a holomorphic g . Then, $f'(0) = g(0)$ and if $|f'(0)| = 1$, then $|g(0)| = 1$, implying that $f(z) = \lambda z$ with $|\lambda| = 1$. \square

Definition. *The automorphism group of a domain Ω is the set of all conformal self-maps of Ω . It is denoted by $\text{Aut}(\Omega)$.*

Remark. The group operation governing $\text{Aut}(\mathbb{D})$ is the composition of maps.

Lemma 3. *Suppose $g \in \text{Aut}(\mathbb{D})$ with $g(0) = 0$. Then $g(z) = e^{i\theta}z$ for $0 \leq \theta \leq 2\pi$.*

Proof. Since $|g(z)| < 1$, by the Schwarz Lemma $|g(z)| \leq |z|$. Moreover, since $g^{-1} \in \text{Aut}(\mathbb{D})$, $|g^{-1}(z)| \leq |z|$ and we conclude that $|g(z)| = |z|$, implying that $g(z) = e^{i\theta}z$ for some $0 \leq \theta \leq 2\pi$. \square

Theorem 4. *Any $f \in \text{Aut}(\mathbb{D})$ is a fractional linear transformation of the form*

$$f(z) = e^{i\theta} \frac{z-a}{1-\bar{a}z} \quad (1)$$

with $a \in \mathbb{C}$ such that $|a| < 1$ and $0 \leq \theta \leq 2\pi$.

Proof. We first show that any map satisfying (1) is in $\text{Aut}(\mathbb{D})$. Let $g(z) = \frac{z-a}{1-\bar{a}z}$, which is a Möbius transform mapping circles and lines to circles and lines. First, note that $g(z)$ maps the unit circle to itself. Indeed, since the complex conjugate has the same modulus, $|e^{i\phi} - a| = |e^{-i\phi} - \bar{a}| = \frac{1}{|e^{i\phi}|} |1 - \bar{a}e^{i\phi}|$ and so

$$|g(e^{i\phi})| = \frac{|e^{i\phi} - a|}{|1 - \bar{a}e^{i\phi}|} = 1.$$

Also, $g(z)$ maps the point $a \in \mathbb{D}$ to the origin since $g(a) = 0$. Crucially, this tells us that f preserves the interior of the unit circle. Thus, if $f(z)$ is as defined above, then $f \in \text{Aut}(\mathbb{D})$.

We now show that any conformal self-map of the unit disk has the form (1). Suppose $h \in \text{Aut}(\mathbb{D})$, so that $a = h^{-1}(0)$ and $(h \circ g^{-1})(0) = h(a) = 0$. By Lemma 2.3, $(h \circ g^{-1})(w) = e^{i\theta}w$ for $0 \leq \theta \leq 2\pi$. Then $h(z) = (h \circ g^{-1})(g(z)) = e^{i\theta}g(z)$ and we have our result. \square

Theorem 5. (Pick's Lemma). *For a holomorphic function $f : \mathbb{D} \rightarrow \mathbb{D}$,*

$$|f'(z)| \leq \frac{1 - |f(z)|^2}{1 - |z|^2}$$

whenever $|z| < 1$, with equality holding if and only if $f \in \text{Aut}(\mathbb{D})$.

Proof. To prove this Lemma, we compose a series of conformal self-maps of \mathbb{D} to map z and $f(z)$ to the origin. Choose $z_0 \in \mathbb{D}$ and let $w_0 = f(z_0)$. Let $g, h \in \text{Aut}(\mathbb{D})$ be such that $g(0) = z_0$ and $h(w_0) = 0$. Suppose $g(z) = \frac{z+z_0}{1+\bar{z}_0z}$ and $h(z) = \frac{w-w_0}{1-\bar{w}_0w}$ so that $g'(0) = 1 - |z_0|^2$ and $h'(w_0) = \frac{1}{1-|w_0|^2}$. Then $(h \circ f \circ g)(0) = 0$ and we can apply Corollary 2 to obtain

$$|(h \circ f \circ g)'(0)| = |h'(w_0)f'(z_0)g'(0)| \leq 1.$$

And therefore

$$|f'(z_0)| \leq \frac{1}{|h'(w_0)||g'(0)|} = \frac{1 - |w_0|^2}{1 - |z_0|^2},$$

as desired.

Now if $f \in \text{Aut}(\mathbb{D})$ then $h \circ f \circ g \in \text{Aut}(\mathbb{D})$ by the closure property of the group. So we obtain equality from Corollary 2 to get that

$$|(h \circ f \circ g)'(0)| = 1.$$

Therefore,

$$|f'(z_0)| = \frac{1 - |w_0|^2}{1 - |z_0|^2}.$$

□

2.1 The hyperbolic metric

Now we define the hyperbolic metric. From above, we know that if $f \in \text{Aut}(\mathbb{D})$ and $w = f(z)$, then

$$\left| \frac{dw}{dz} \right| = \frac{1 - |w|^2}{1 - |z|^2} \Rightarrow \frac{|dw|}{1 - |w|^2} = \frac{|dz|}{1 - |z|^2}.$$

Suppose that γ is a smooth curve in \mathbb{D} . Then,

$$\int_{f \circ \gamma} \frac{|dw|}{1 - |w|^2} = \int_{\gamma} \frac{|dz|}{1 - |z|^2}. \tag{2}$$

This helps us obtain a metric that is invariant under mappings in $\text{Aut}(\mathbb{D})$, leading us to the hyperbolic metric.

Definition. The *hyperbolic length* of γ is defined as $2 \int_{\gamma} \frac{|dz|}{1 - |z|^2}$.

Here, the coefficient 2 does not matter. We choose it only so that the curvature of the metric is adjusted to -1 .

Definition. The *hyperbolic distance* between $z_0, z_1 \in \mathbb{D}$ is the infimum of the hyperbolic lengths of all piecewise smooth curves γ from z_0 to z_1 in \mathbb{D} . This is denoted by $\rho(z_0, z_1)$.

By (2.2), since the hyperbolic length of a piecewise smooth curve γ from z_0 to z_1 in \mathbb{D} is invariant under conformal self-maps of \mathbb{D} , so is the hyperbolic distance between z_0 and z_2 . In particular, for any $f \in \text{Aut}(\mathbb{D})$,

$$\rho(f(z_0), f(z_1)) = \rho(z_0, z_1)$$

for any $z_0, z_1 \in \mathbb{D}$. This leads us to the concept of hyperbolic geodesics.




Figure 1: Geodesics on \mathbb{D}

Theorem 6. The geodesic, or the path of shortest hyperbolic length between two points z_0 and z_1 is given by the arc of the circle that passes through them and intersects the boundary of \mathbb{D} orthogonally. This is illustrated in Fig. 1.

Proof. Suppose $f \in \text{Aut}(\mathbb{D})$ such that $f(z_0) = 0$. By multiplying it by a constant $e^{i\tau}$, we can rotate f such that $f(z_1) = r > 0$, i.e. $f(z_1)$ lies on the positive real axis. Since f is a conformal mapping, any circle orthogonal to \mathbb{D} will be mapped to a circle orthogonal to the boundary of \mathbb{D} . Therefore, we need only to show that the unique shortest length path from $f(z_0)$ to $f(z_1)$ is the straight line segment from 0 to r along the real axis pictured in Fig. 2.

Parametrize the path $\gamma(t) = x(t) + iy(t)$, $0 \leq t \leq 1$ from 0 to r , taking values in \mathbb{D} . The path $\alpha(t) = \text{Re}(\gamma(t)) = x(t)$ gives the path along the real axis, and the result follows from the fact that $|x(t)| \leq |\gamma(t)|$. Note that

$$\int_{\alpha} \frac{|dz|}{1 - |z|^2} = \int_0^1 \frac{|x'(t)dt|}{1 - x(t)^2} \leq \int_0^1 \frac{|\gamma'(t)dt|}{1 - |\gamma(t)|^2} = \int_{\gamma} \frac{|dz|}{1 - |z|^2}$$

If $y(t) \neq 0$ for some t , then $|x(t)| < |\gamma(t)|$ and the above inequality is strict, telling us that the shortest length path must be along the x -axis. If, suppose, the parametrized path were decreasing on some interval of the real axis, then the parametrization $\alpha(t)$ would not be one-to-one. In that case, we remove the interval of t 's along which $\alpha(t)$ starts and ends at the same value to obtain a shorter path, contradicting minimality. Therefore, the shortest length path must be the real, nondecreasing line segment from 0 to r . □




Figure 2: The unique shortest length path from $f(z_0)$ to $f(z_1)$.

Theorem 7. Every holomorphic function $f : \mathbb{D} \rightarrow \mathbb{D}$ is a contraction mapping under the hyperbolic metric. In other words,

$$\rho(f(z_0), f(z_1)) \leq \rho(z_0, z_1)$$

for any $z_0, z_1 \in \mathbb{D}$. If $f \in \text{Aut}(\mathbb{D})$, then equality holds. Otherwise, $\rho(f(z_0), f(z_1)) < \rho(z_0, z_1)$ for all $z_0, z_1 \in \mathbb{D}$, $z_0 \neq z_1$.

Proof. Suppose γ is the geodesic from z_0 to z_1 . Then, $f \circ \gamma$ parametrizes the curve from $w_0 = f(z_0)$ to $w_1 = f(z_1)$. The hyperbolic distance is the infimum of the hyperbolic lengths, so we use this and apply Pick's Lemma to obtain

$$\begin{aligned} \rho(f(z_0), f(z_1)) &\leq 2 \int_{f \circ \gamma} \frac{|dw|}{1-|w|^2} = 2 \int_{\gamma} \frac{|f'(z)|}{1-|f(z)|^2} \\ &\leq 2 \int_{\gamma} \frac{|dz|}{1-|z|^2} = \rho(z_0, z_1) \end{aligned}$$

Where the last inequality comes from Pick's Lemma. In the case where $f \notin \text{Aut}(\mathbb{D})$, strict equality holds from the statement of Pick's Lemma. \square

For any point ζ_0 on the boundary $\partial\mathbb{D} = \{z \in \mathbb{C} : |z| = 1\}$, the hyperbolic distance from 0 to ζ_0 must be infinite. To see this, we take the integral as usual

$$\begin{aligned} \rho(0, \zeta_0) &= \int_0^{|\zeta_0|} \frac{dz}{1-z^2} = \log \left(\frac{1+z}{1-z} \right) \Bigg|_0^{z=|\zeta_0|} \\ &= \log \left(\frac{1-|\zeta_0|}{1+|\zeta_0|} \right) = \infty. \end{aligned}$$

The concept of infinity is thus well defined in terms of spherical distance. For the extension of the Riemann mapping theorem to the uniformization theorem, it is useful to note that there is also a direct correspondence between the Euclidean metric and the metric induced on the Riemann sphere, or the extended complex plane \mathbb{C}^* . This prompts us to introduce the metric on the Riemann sphere.

2.2 The spherical metric

Recall that the stereographic projection of a point P on the Riemann sphere is the point $z = x + iy$ where the line through the north pole $N = (0, 0, 1)$ and $P = (X, Y, Z)$ in Euclidean space meets the plane $Z = 0$. We parametrize this line in the usual way as $L(t) = N + t(P - N)$, $-\infty < t < \infty$. To find the coordinates of its projection, i.e. the point $(x, y, 0)$ where $L(t)$ meets the x, y -plane, we solve for the following

$$\begin{aligned} (x, y, 0) &= (0, 0, 1) + t[(X, Y, Z) - (0, 0, 1)] \\ &= (tX, tY, tZ - t) \end{aligned}$$

The chordal metric induced on \mathbb{C} is defined in terms of this line and the chordal distance between two point z and w is given by

$$\frac{2|z-w|}{\sqrt{1+|z|^2}\sqrt{1+|w|^2}}.$$

The infinitesimal version of this is $\frac{2|dz|}{1+|z|^2}$, and so the length of a curve in the spherical metric is found by integrating this over the curve on the Riemann sphere.

Definition. Suppose that $\gamma(t)$ parametrizes a path from z_0 to z_1 in $\mathbb{C}^* = \mathbb{C} \cup \{\infty\}$. The spherical length of γ is defined as $2 \int_{\gamma} \frac{|dz|}{1+|z|^2} = 2 \int \frac{|\gamma'(t)|}{1+|\gamma(t)|^2} dt$.

Definition. The spherical distance between $z_0, z_1 \in \mathbb{C}^*$ is the infimum of the spherical lengths of all paths from z_0 to z_1 . This is denoted by $\sigma(z_0, z_1)$ and is invariant under rotation mappings.

3 NORMAL FAMILIES OF FUNCTIONS

To establish the existence of the Riemann mapping, we require a compactness principle for families of holomorphic functions. This is provided by the theory of normal families, which characterizes collections of functions with convergent subsequences. In Section 4, we will employ this theory to introduce the theorems of Montel and Hurwitz, which are fundamental to our proof.

Definition. A sequence of holomorphic functions $\{f_n(z)\}$ on a domain $\Omega \subseteq \mathbb{C}$ is said to **converge normally** to a holomorphic function $f(z)$ on Ω if it converges uniformly on each compact subset of Ω . Under the spherical metric, this would imply that $\sigma(f_n(z), f(z)) \xrightarrow{\text{unif.}} 0$ on compact subsets of Ω as $n \rightarrow \infty$.

Definition. A family \mathcal{F} of holomorphic functions on a domain $\Omega \subseteq \mathbb{C}$ is called a **normal family** if every sequence in \mathcal{F} has a subsequence that converges normally in Ω .

Theorem 8. (Hurwitz). Let $\{f_n(z)\}$ be a sequence of holomorphic functions on a domain Ω that converges normally in Ω to a holomorphic function $f(z)$ and suppose $f(z)$ has a zero of order N at $z_0 \in \mathbb{C}$. Then there exist $\rho, K > 0$ such that for all $n > K$, $f_n(z)$ has exactly N zeroes, counted with multiplicity, in the disk $\{|z - z_0| < \rho\}$ and these zeroes converge to z_0 as $n \rightarrow \infty$.

Proof. Since $\Omega \subseteq \mathbb{C}$ is open, we can choose $\rho > 0$ small enough so that $\{|z - z_0| \leq \rho\} \subseteq \Omega$ and this disk contains no other zeroes of $f(z)$, i.e. $f(z) \neq 0$ for $0 < |z - z_0| \leq \rho$. Therefore, there exists $\delta > 0$ such that $|f(z)| \geq \delta$ on the circle $|z - z_0| = \rho$.

For such a δ , since $f_n(z) \xrightarrow{\text{unif.}} f(z)$ for $|z - z_0| \leq \rho$, for n large we have that $|f_n(z)| > \frac{\delta}{2}$ on the circle $C_\rho = \{|z - z_0| = \rho\}$. Moreover, $f'_n(z) \xrightarrow{\text{unif.}} f'(z)$ by the Cauchy integral formula, and thus $\frac{f'_n(z)}{f_n(z)} \xrightarrow{\text{unif.}} \frac{f'(z)}{f(z)}$ on C_ρ . Since all $f_n(z)$ and $f(z)$ are holomorphic, by the Argument Principle

$$\lim_{n \rightarrow \infty} N_0(f_n) = \lim_{n \rightarrow \infty} \frac{1}{2\pi i} \int_{C_\rho} \frac{f'_n(z)}{f_n(z)} dz = \frac{1}{2\pi i} \int_{C_\rho} \frac{f'(z)}{f(z)} dz = N$$

where N_0 denotes the number of zeroes of the sequence inside the disk $\{|z - z_0| < \rho\}$, counted with multiplicity. For a large enough n , $N_0(f_n) = N$ and we are done. \square

Definition. A function f is **univalent** in the domain Ω if it is holomorphic and injective on Ω .

Theorem 9. If $f_n(z)$ is a sequence of univalent functions on a domain Ω converging normally on Ω to a function $f(z)$, then $f(z)$ is either univalent or constant.

Proof. Assume that $f(z)$ is not constant. Suppose that $f(z_0) = f(\zeta_0) = w_0$ for $z_0, \zeta_0 \in \Omega$. Then, z_0 and ζ_0 are zeroes of the function $f(z) - w_0$ and by Theorem 3.1, there are sequences $z_n \rightarrow z_0$ and $\zeta_n \rightarrow \zeta_0$ such that $f_n(z_n) - w_0 = 0$ and $f_n(\zeta_n) - w_0 = 0$. That is, z_n and ζ_n are zeroes of the function $f_n(z) - w_0$. Since the f_n 's are univalent, we have $z_n = \zeta_n$ and in the limit $z_0 = \zeta_0$. \square

4 MONTEL'S THEOREM AND EXTREMAL PROBLEMS

The construction of the Riemann map reduces to solving an extremal problem in function theory. We apply Montel's Theorem to obtain a candidate mapping that maximizes the derivative at a fixed point. This section establishes the existence of extremal functions through compactness arguments, a crucial step that will enable us to prove the Riemann map's uniqueness in Section 5.

Definition. Suppose $A \subseteq \mathbb{C}$ is a subset of the complex plane. A family \mathcal{F} of functions is **equicontinuous** (in the spherical metric) at a point $z_0 \in A$ if for any $\varepsilon > 0$, there exists a $\delta > 0$ such that if $z \in A$ satisfies $|z - z_0| < \delta$ then $\sigma(f(z) - f(z_0)) < \varepsilon$ for all $f \in \mathcal{F}$.

¹Recall that the Bolzano-Weierstrass theorem states that any bounded sequence of complex numbers has a convergent subsequence.

Definition. The family \mathcal{F} is **uniformly bounded** on A if there exists an $M > 0$ such that $|f(z)| \leq M$ for all $f \in \mathcal{F}$.

Definition. A subset $A \subseteq \mathbb{C}$ is **compact** if it is closed and bounded.

We are now prepared to state a version of the Arzelà-Ascoli theorem, an important measure-theoretic result that will allow us to state Montel's theorem.

Theorem 10. Let $A \subseteq \mathbb{C}$ be compact and \mathcal{F} a uniformly bounded family of continuous functions on A . If \mathcal{F} is equicontinuous at every point of A with respect to the spherical metric, then each sequence of functions in \mathcal{F} has a subsequence that converges normally on A in the spherical metric.

Proof. Let $\{f_n(z)\}$ be a sequence of functions in \mathcal{F} and S be a set of points that is dense in A . Since \mathcal{F} is uniformly bounded, the set $\{f(z_1)\}$ is bounded for $z_1 \in S$, and by the Bolzano-Weierstrass theorem¹, there is a sequence of functions $\{f_{n_1}\} \in \mathcal{F}$ such that $\{f_{n_1}(z_1)\}$ converges. Moreover, there is a subsequence $\{f_{n_2}\}$ of the bounded sequence $\{f_{n_1}(z_2)\}$ such that $\{f_{n_2}(z_2)\}$ converges. Repeating this diagonalization argument for the nested subsequences

$$f_{n_1} \supseteq f_{n_2} \supseteq \dots$$

in \mathcal{F} , we find a subsequence $\{f_{n_k}\}$ that converges at z_1, \dots, z_k . From this, we construct the diagonal subsequence $\{f\}$ whose m th term is the m th term in $\{f_{n_m}\}$, so that $\{f\}$ converges at every point in S .

By convergence, for any $\varepsilon > 0$ and $z_k \in S$, there is an integer $N(\varepsilon, z_k)$ such that for all $n, m \geq N(\varepsilon, z_k)$, $\sigma(f_m(z_k), f_n(z_k)) < \frac{\varepsilon}{3}$. Moreover, since \mathcal{F} is equicontinuous, for every $z \in A$ there is a subdomain U_z of the complex plane containing z such that $\sigma(f(\zeta), f(\xi)) < \frac{\varepsilon}{3}$ for all $f \in \mathcal{F}$ and $\zeta, \xi \in U_z \cap S$.

The collection of subsets $\{U_z\}_{z \in A}$ forms an open cover of A . Since $A \subseteq \mathbb{C}$ is compact, this admits a finite subcover U_1, \dots, U_J . Moreover, there is some integer K such that each U_j , $1 \leq j \leq J$, contains a point $z_k \in S$ for $1 \leq k \leq K$. Moreover for any $\xi \in A$, there are j and k such that $\xi, z_k \in U_j$. For this k ,

$$\begin{aligned} \sigma(f_m(\xi), f_n(\xi)) &\leq \sigma(f_n(\xi), f_n(z_k)) + \sigma(f_n(z_k), f_m(z_k)) \\ &\quad + \sigma(f_m(z_k), f_m(\xi)) \\ &< \frac{\varepsilon}{3} + \frac{\varepsilon}{3} + \frac{\varepsilon}{3} = \varepsilon. \end{aligned}$$

for $n, m \geq \max(N(\varepsilon, z_1), \dots, N(\varepsilon, z_K))$. Therefore, the sequence is uniformly Cauchy in A and converges normally to a continuous function. \square

Next, we state the thesis-grade Montel's theorem with an elementary proof provided by Gamelin [1]. One of the intermediate steps is proving that a family of functions is bounded if the derivatives of its functions are uniformly bounded. We formalize it in the following result.

Theorem 11. *Suppose that \mathcal{F} is a family of functions on Ω and there is an $M > 0$ such that $|f'(z)| < M$ for all $z \in \Omega$ and all $f \in \mathcal{F}$. Then, \mathcal{F} is equicontinuous at every point $z_0 \in \Omega$.*

Proof. Suppose $z_0 \in \Omega$ and $\varepsilon > 0$. Set $\delta = \frac{\varepsilon}{M}$. Then, for any $z \in \Omega$ such that $|z - z_0| < \delta$ we integrate along the line segment between z and z_0 to get

$$|f(z) - f(z_0)| = \left| \int_{z_0}^z f'(\zeta) d\zeta \right| \leq M|z - z_0| < M \cdot \frac{\varepsilon}{M} = \varepsilon.$$

Since δ only depends on ε and M , \mathcal{F} is equicontinuous at z_0 . \square

Theorem 12 (Montel). *If \mathcal{F} is a family of holomorphic functions on Ω such that \mathcal{F} is uniformly bounded on compact subsets of Ω , then each sequence in \mathcal{F} has a subsequence that converges normally in Ω to a holomorphic function $f(z)$.*

Proof. Since Ω is open, if $z_0 \in \Omega$ then there is an $r > 0$ such that the disk $C_r = \{|z - z_0| \leq r\}$ is in Ω and \mathcal{F} is uniformly bounded on C_r . By the Cauchy integral formula, the derivatives of f are uniformly bounded on the smaller disk $C'_r = \{|z - z_0| \leq r - \varepsilon\}$. Therefore, by Theorem 11, \mathcal{F} is equicontinuous at every point $z_0 \in \Omega$.

Now, let A_n be the set of $z \in \Omega$ such that $|z| \leq n$ and the distance from z to $\partial\Omega$ is at least $\frac{1}{n}$. Then the sequence $\{A_n\}_{n=1}^\infty$ has the following properties: each A_n is compact, increases to Ω , and each compact subset of Ω is contained in some A_n . By Arzelà-Ascoli, if $\{f_n\}$ is a sequence in \mathcal{F} , it has a subsequence f_{11}, f_{12}, \dots that converges uniformly on A_1 . This has a further subsequence f_{21}, f_{22}, \dots that converges on A_2 and so on. Then the diagonal subsequence f_{11}, f_{22}, \dots converges uniformly on each A_n and therefore normally on Ω . \square

The remaining proof requires an introduction to the notion of "extremal problems". In 1947, Lars Ahlfors studied [2] the family \mathcal{F} of holomorphic functions $f(z)$ with $|f(z)| \leq 1$ and considered the extremal problem that maximizes the derivative $|f'(z_0)|$ across all $f \in \mathcal{F}$ at a point $z_0 \in \Omega$.

4.1 Extremal problems

Definition. *For a given $z_0 \in \Omega$, the extremal value for Ahlfors' problem is*

$$C = \sup_{f \in \mathcal{F}} |f'(z_0)|.$$

Definition. *An extremal function for Ahlfors' problem is a function $G \in \mathcal{F}$ such that $|G'(z_0)| = C$.*

Montel's theorem guarantees the existence of an extremal function as follows:

Theorem 13. *Suppose Ω is a domain in \mathbb{C} with a non-constant bounded holomorphic function on it, and $z_0 \in \Omega$. Then there exists a holomorphic function $G(z)$ on Ω such that $|G(z)| \leq 1$ for all $z \in \Omega$ and $|f'(z_0)| \leq |G'(z_0)|$ for any holomorphic f on Ω such that $|f(z)| \leq 1$ for all $z \in \Omega$. Moreover, $G(z_0) = 0$ and $G'(z_0) \neq 0$.*

Proof. Let \mathcal{F} be the family of functions f on Ω such that $|f(z)| \leq 1$ for all $z \in \Omega$ and $f(z_0) = 0$. Suppose $\{f_n\}$ is a sequence of functions in \mathcal{F} such that $|f'_n(z_0)|$ converges to C . By Montel's theorem, there is a subsequence that converges normally on Ω to a holomorphic function $G(z)$ such that $|G(z)| \leq 1$ and $|G'(z_0)| = C$.

To prove the final statement, suppose that $h(z)$ is a bounded nonconstant holomorphic function on Ω . Then $h(z) - h(z_0)$ has a zero of order N at z_0 . For $\varepsilon > 0$, consider the function

$$f(z) = \frac{\varepsilon(h(z) - h(z_0))}{(z - z_0)^{N-1}}.$$

For a small enough ε , $f(z)$ satisfies $|f(z)| \leq 1$ and $f'(z_0) \neq 0$. Therefore, $C = |G'(z_0)| > 0$. Also, the function $g(z) = \frac{G(z) - G(z_0)}{1 - \overline{G(z_0)}z}$ is the composition of $G(z)$ and a conformal self-map of \mathbb{D} , therefore $g \in \mathcal{F}$. Finally,

$$|G'(z_0)| \geq |g'(z_0)| = \frac{|G'(z_0)|}{1 - |G(z_0)|^2},$$

so we have $|G(z_0)| = 0$, as desired. \square

5 PROOF OF THE RIEMANN MAPPING THEOREM

Now, it remains to show that if Ω is any simply connected domain such that $\Omega \neq \mathbb{C}$, then it is conformally equivalent to \mathbb{D} . For the final part of the proof, we exploit the mapping properties of the square root function $\sqrt{z} = e^{\frac{1}{2} \log(z)}$.

Lemma 14. *Suppose that Ω is a simply connected domain and $a \notin \Omega$. Let $h(z)$ be a holomorphic branch of $\sqrt{z - a}$ in Ω . Then $h(\Omega)$ is univalent on Ω and $h(\Omega)$ is disjoint from $-h(\Omega)$.*

Proof. Suppose that $\Omega \subsetneq \mathbb{C}$ is a simply connected domain and $a \in \mathbb{C} \setminus \Omega$. By simple connectivity, there is a holomorphic branch $g(z)$ of $\log(z-a)$ in Ω . Then, we take the holomorphic branch of $\sqrt{z-a}$ defined by $h(z) = e^{\frac{g(z)}{2}}$. To show that h is univalent, note that $h(z)^2 = z-a \neq 0$ in Ω . Suppose that $h(z_1) = h(z_2)$ for some $z_1, z_2 \in \Omega$. Then, $z_1 = h(z_1)^2 + a = h(z_2)^2 + a = z_2$ in Ω .

We also show that if $w_0 \in h(\Omega)$, then $-w_0 \notin h(\Omega)$. For contradiction, suppose $w_0 = h(z_0)$ and $-w_0 = h(z_1)$ for some $z_0, z_1 \in \Omega$. Then, $z_0 = w_0^2 + a = w_1^2 + a = z_1$, which is only possible when $w_0 = w_1 = 0$. However, this cannot occur since $0 \notin h(\Omega)$. \square

Corollary 15. *If the disk $C_\varepsilon = \{w : |w - w_0| \leq \varepsilon\} \subseteq h(\Omega)$, then its negative is disjoint from $h(\Omega)$ and therefore $|h(z) - (-w_0)| > \varepsilon$ for $z \in \Omega$. This implies that for all $z \in \Omega$, $\mu(z) = \frac{\varepsilon}{|h(z)+w_0|} < 1$, and therefore μ holomorphically maps Ω to a subdomain of \mathbb{D} .*

Let $\rho(\zeta, \xi)$ be the hyperbolic metric in \mathbb{D} . From Theorem 2.7, the globally holomorphic square mapping ζ^2 from \mathbb{D} onto itself is a contraction in the hyperbolic metric. Moreover, for any $\varepsilon > 0$, ζ^2 is a strict contraction on the disk $\{|\zeta| \leq 1 - \varepsilon\}$, i.e. there exists a $c > 0$ such that $\rho(\zeta^2, \xi^2) \leq c\rho(\zeta, \xi)$. Setting $\sqrt{z} = \zeta$, we see that the square-root mapping is expanding on subsets of \mathbb{D} with respect to the hyperbolic metric. Analogous to the square mapping case, for any $0 < r < 1$, there is a $C > 1$ such that

$$\rho(\pm\sqrt{\zeta}, \pm\sqrt{\xi}) \geq C\rho(\zeta, \xi)$$

whenever $|\zeta|, |\xi| \leq r$.

Lemma 16. *Suppose that Ω is a simply connected domain $\Omega \subset \mathbb{D}$ with $0 \in \Omega$. If $\Omega \neq \mathbb{D}$, then there exists a conformal map $\psi(z)$ onto a subdomain of \mathbb{D} such that $\psi(0) = 0$ and $|\psi'(0)| > 1$.*

Proof. Let $g \in \text{Aut}(\mathbb{D})$ be such that there is a point $a \in \mathbb{D} \setminus \Omega$ with $g(a) = 0$. By the bijectivity of g , the image $g(\Omega)$ must be a simply connected subdomain of \mathbb{D} not containing 0. Then, by the previous Lemma there is a holomorphic branch $h(\xi)$ of $\sqrt{\xi}$ that is univalent on $g(\Omega)$ and $h \circ g$ is an expanding map from \mathbb{D} to itself. Let $f \in \text{Aut}(\mathbb{D})$ be such that $f(h(g(0))) = 0$, and set $\psi = f \circ h \circ g$. Then $\psi : \Omega \rightarrow A \subseteq \mathbb{D}$ is a conformal map onto a subdomain A of \mathbb{D} such that $\psi(0) = 0$. Since $f, g \in \text{Aut}(\mathbb{D})$, they preserve hyperbolic distances whereas h strictly expands them, implying that ψ is a strict expansion near 0. Thus, for each $\varepsilon > 0$, there is a $C > 1$ such that

$$\frac{\rho(\psi(\zeta), 0)}{\rho(\zeta, 0)} \geq C$$

for $|\zeta| < \varepsilon$. From the definition of the hyperbolic metric, as $\zeta \rightarrow 0$, $\rho(\zeta, 0) \approx 2|\zeta|$ and $\rho(\psi(\zeta), 0) \approx 2|\psi(\zeta)|$, giving

$$|\psi'(0)| = \lim_{\varepsilon \rightarrow 0} \frac{|\psi(\zeta)|}{|\zeta|} \approx \frac{\rho(\psi(\zeta), 0)}{\rho(\zeta, 0)} \geq C$$

and therefore $|\psi'(0)| > 1$. \square

Now, we restate the Riemann mapping theorem and complete the proof using Lemmas 5.1 and 5.3.

Theorem 17. (Riemann Mapping Theorem). *Suppose that $\Omega \subset \mathbb{C}$ is a simply connected domain such that $\Omega \neq \mathbb{C}$. Then, there exists a conformal map, called the Riemann map, of Ω onto \mathbb{D} .*

Proof. For the final proof, suppose that $\Omega \neq \mathbb{C}$ is a simply connected domain and that for some fixed $z_0 \in \Omega$, \mathcal{F} is the family of univalent functions $f(z)$ on Ω such that $|f(z)| < 1$ everywhere and $f(z_0) = 0$. Since z_0 is a zero of order 1, for any bounded conformal map $\mu(z)$ from Ω such as that defined in Corollary 15, there is some $\varepsilon > 0$ such that the map $f(z) = \varepsilon(\mu(z) - \mu(z_0))$ satisfies $|f(z)| < 1$ and $f(z_0) = 0$, so $f \in \mathcal{F}$ and \mathcal{F} is nonempty. We attempt to solve the extremal problem of maximizing $|f'(z_0)|$ over \mathcal{F} . As in Ahlfors' problem, set

$$C = \sup_{f \in \mathcal{F}} |f'(z_0)| > 0$$

and let $\{f_n\} \subseteq \mathcal{F}$ be a sequence of functions such that $|f'_n(z_0)| \rightarrow C$. We formulate the Riemann map as the extremal function for this problem. By Montel's theorem there exists a subsequence of $\{f_n\}$ that converges normally on Ω to a holomorphic function $\phi(z)$. Moreover $|\phi(z)| \leq 1$, $\phi(z_0) = 0$ and $|\phi'(z_0)| = C$. By Theorem 3.2, ϕ is either constant or univalent but since $\phi'(z_0) \neq 0$, ϕ must be univalent from Ω to a subdomain of \mathbb{D} .

It must be the case that $\phi(\Omega) = \mathbb{D}$. Otherwise, $0 \in \phi(\mathbb{D})$ and by Lemma 16 we could compose ϕ with a conformal map $\psi : \phi(\Omega) \subsetneq \mathbb{D} \rightarrow \mathbb{D}$ such that $\psi(0) = 0$ and $|\psi'(0)| > 1$. Then, the function $\psi \circ \phi \in \mathcal{F}$ would satisfy $|(\psi \circ \phi)'(z_0)| = |\psi'(0)\phi'(z_0)| > C$, contradicting the definition of C . \square

REFERENCES

- [1] Gamelin, T. Complex Analysis. (Springer New York, 2001)
- [2] Ahlfors, L. Bounded analytic functions. *Duke Mathematical Journal*. **14**, 1-11 (1947)

THE MINIMUM-NORM OLS ESTIMATOR: UNBIASEDNESS AND OPTIMALITY IN THE OVERPARAMETERIZED REGIME

Margaux Butterfield, Perry Santry, Noé Vartanian

This paper investigates the behavior of the Ordinary Least Squares (OLS) estimator as the dimensionality of the covariates, p , crosses the sample size threshold, n . We analyze the transition from the classical regime ($p < n$) where the OLS solution is unique, to the overparameterized regime ($p \geq n$), where the model interpolates the data and admits infinitely many solutions. Focusing on the minimum ℓ_2 -norm estimator defined by the Moore-Penrose pseudoinverse, $\hat{\beta} = X^\dagger y$, we establish a high-dimensional analogue of the Gauss-Markov theorem. We demonstrate that the minimum-norm estimator is the strictly unique linear estimator that is unbiased for the identifiable component of the true parameter, $\beta^* = X^\dagger X\beta$. Comparing it against the broader class of minimum-residual (interpolating) estimators, theoretical results prove that all interpolating estimators share identical variance over the row space of the design matrix, but the minimum-norm solution is optimal by maintaining strictly zero variance in the null space. These findings are corroborated by simulations which explore the behaviour of the minimum-norm estimator in a stochastic model.

1 INTRODUCTION

Linear regression is a cornerstone of statistical learning, traditionally understood through the lens of the bias-variance trade-off. In the classical regime, where the number of observations n exceeds the number of parameters p , the Ordinary Least Squares (OLS) estimator is unique and well-behaved. However, modern applications in machine learning frequently operate in the overparameterized regime ($p \geq n$), where the model complexity allows for perfect interpolation of training data. This shift challenges conventional statistical intuition, as achieving zero training error usually indicates severe overfitting and poor generalization in classical statistics.

The goal of this paper is to mathematically formalize the behavior of the OLS solution as it transitions across these regimes. We begin by reviewing the properties of the OLS estimator in the classical setting, utilizing the standard normal equations. We then extend this analysis to the interpolation threshold and beyond, where the design matrix X becomes full row rank but rank-deficient in columns, leading to an underdetermined system.

To resolve the non-uniqueness of solutions in the overparameterized regime, we introduce the minimum ℓ_2 -norm estimator, $\hat{\beta}$, derived via the Moore-Penrose pseudoinverse. A central contribution of this work is the adaptation of the Gauss-Markov theorem to high-dimensional settings. We define an appropriate notion of unbiasedness for the projected parameter β^* and prove that the minimum-norm estimator is the unique linear row-space-unbiased estimator. We then establish its optimality against the broader class of minimum-residual (interpolating) estimators. Finally, we validate our theoretical derivations through a series of simulations, confirming that while all interpolating estimators capture the signal equally well, the minimum-norm solution uniquely avoids variance leakage into the null

space.

2 THE OLS ESTIMATOR ACROSS REGIMES

Throughout this paper, we work with an observed dataset

$$D_n = \{(\mathbf{x}_i, y_i)\}_{i=1}^n,$$

where each $\mathbf{x}_i \in \mathbb{R}^p$ represents a vector of covariates and $y_i \in \mathbb{R}$ is its associated response. We gather the covariates into the design matrix $X \in \mathbb{R}^{n \times p}$ and the responses into the vector $\mathbf{y} \in \mathbb{R}^n$.

We assume throughout that the design matrix X has full rank, that is:

$$\text{rank}(X) = \min\{n, p\}.$$

We work under the standard linear regression model

$$y_i = \mathbf{x}_i^\top \beta + \varepsilon_i, \quad i = 1, \dots, n.$$

We examine how the behaviour of the least-squares estimator changes as the number of parameters p crosses the sample size n .

2.1 The Ordinary Least Square Estimator (OLS)

Definition (Ordinary least squares loss). *Given a dataset $D_n = \{(x_i, y_i)\}_{i=1}^n$, the **ordinary least squares (OLS) loss** associated with $\beta \in \mathbb{R}^p$ is defined as*

$$L(\beta) = \frac{1}{n} \|\mathbf{y} - X\beta\|_2^2.$$

This quantity measures the average squared difference between the observed responses and the linear predictions generated by β .

Definition (Ordinary least squares estimator). *The ordinary least squares estimator is any vector $\hat{\beta} \in \mathbb{R}^p$ that minimizes the OLS loss, that is,*

$$\hat{\beta} \in M,$$

where

$$M := \arg \min_{\beta \in \mathbb{R}^p} L(\beta).$$

Definition (Predicted Value). *Given a design matrix X and an estimator $\hat{\beta}$, the corresponding vector of predicted responses is*

$$\hat{y} = X\hat{\beta}.$$

Theorem 1. *$\hat{\beta}$ satisfies the normal equations, that is:*

$$(X^\top X)\hat{\beta} = X^\top \mathbf{y}.$$

Note that no assumptions have to be made for Theorem 2.

2.2 The Classical Regime, $p < n$

Under the assumption that X has full column rank, $X^\top X$ is invertible and thus:

Theorem 2. *If X has full column rank, then $\hat{\beta}$ is unique and satisfies:*

$$\hat{\beta} = (X^\top X)^{-1} X^\top \mathbf{y}.$$

2.3 The Interpolation Threshold, $p = n$ and Overparameterization, $p > n$

In the following section, we will consider the case where the number of covariates is equal to or exceeds the sample size ($p \geq n$). Under these conditions, the model is called overparameterized.

Theorem 3. *If the design matrix X has full row rank then:*

$$\hat{\mathbf{y}} = \mathbf{y}$$

i.e., the model *interpolates* the data.

Proof. Since $\hat{\beta}$ is an OLS estimator, it satisfies the normal equations by Theorem 2:

$$\begin{aligned} X^\top X \hat{\beta} &= X^\top \mathbf{y} \\ \implies X^\top \hat{\mathbf{y}} &= X^\top \mathbf{y} && \text{(Since } \hat{\mathbf{y}} = X\hat{\beta}\text{)} \\ \implies X^\top (\hat{\mathbf{y}} - \mathbf{y}) &= \mathbf{0} \\ \implies \hat{\mathbf{y}} - \mathbf{y} &\in \ker(X^\top). \end{aligned}$$

Since X has full row rank, $\text{rank}(X) = n$. Thus:

$$\text{rank}(X^\top) = \text{rank}(X) = n,$$

so by the Rank–Nullity theorem,

$$\dim \ker(X^\top) = n - \text{rank}(X^\top) = 0,$$

which implies:

$$\ker(X^\top) = \{\mathbf{0}\}.$$

Therefore,

$$\begin{aligned} \hat{\mathbf{y}} - \mathbf{y} &= \mathbf{0} \\ \implies \hat{\mathbf{y}} &= \mathbf{y}. \end{aligned}$$

□

Thus, when $p \geq n$, the OLS estimator achieves zero training error. The model has no bias on the training sample D_n , but because it interpolates the data, its variance becomes extremely large.

The case $p = n$ is the interpolation threshold. In this situation, $\hat{\beta}$ is unique, and the model fits the data exactly. On the other hand, when $p > n$, the system $X\beta = y$ becomes underdetermined, which leads to infinitely many interpolating solutions for β .

Minimum-Norm Solution

Having in mind that, in the overparameterized regime multiple models minimize the OLS loss, it seems important to establish a rule that would allow us to find a unique $\hat{\beta}$. The standard convention is to consider the minimum l_2 -norm ordinary least squares estimator. That is, among all OLS estimators, the one that satisfies

$$\hat{\beta}_{\text{norm}} \in \arg \min_{\beta \in M} \|\beta\|_2.$$

Let us now introduce the concept of the pseudoinverse.

Definition (Moore–Penrose pseudoinverse). *For a matrix $X \in \mathbb{R}^{n \times p}$, the Moore–Penrose pseudoinverse $X^\dagger \in \mathbb{R}^{p \times n}$ is the unique matrix satisfying the following four conditions:*

1. $XX^\dagger X = X$,
2. $X^\dagger XX^\dagger = X^\dagger$,
3. $(XX^\dagger)^\top = XX^\dagger$,
4. $(X^\dagger X)^\top = X^\dagger X$.

Note that the Moore–Penrose pseudoinverse X^\dagger is guaranteed to exist (constructed via the SVD) and is unique for each matrix $X \in \mathbb{R}^{n \times p}$. Also note that its properties imply that XX^\dagger and $X^\dagger X$ are orthogonal projections onto the column space and row space of X , respectively [3].

As a consequence, the minimum l_2 -norm OLS estimator is exactly $\hat{\beta}_{\text{norm}} = X^\dagger \mathbf{y}$. Before proving this, we first prove a necessary result.

Theorem 4 (Form of OLS estimators). *If $\hat{\beta}$ minimises the OLS loss $\|X\hat{\beta} - y\|_2^2$, then*

$$\hat{\beta} = X^\dagger y + (I_p - X^\dagger X)w,$$

for $w \in \mathbb{R}^p$.

Proof. Since $\hat{\beta}$ is an OLS estimator, it satisfies the normal equations by Theorem 2:

$$X^\top X \hat{\beta} = X^\top y.$$

Multiplying by $(X^\dagger)^\top$ on the left gives:

$$\begin{aligned} (X^\dagger)^\top X^\top X \hat{\beta} &= (X^\dagger)^\top X^\top y \\ \implies (XX^\dagger)^\top X \hat{\beta} &= (XX^\dagger)^\top y && ((X^\dagger)^\top X^\top = (XX^\dagger)^\top) \\ \implies XX^\dagger X \hat{\beta} &= XX^\dagger y && ((XX^\dagger)^\top = XX^\dagger) \\ \implies X \hat{\beta} &= XX^\dagger y. && (XX^\dagger X = X) \end{aligned}$$

Thus $\hat{\beta}$ solves $X \hat{\beta} = XX^\dagger y$ (*).

Let $\beta_0 = X^\dagger y$, then β_0 satisfies (*). By Proposition 2.5 of [5]. If a system of the form $Ax = b$ has a particular solution p , then any other solution has the form $p + x$, where x is an arbitrary element of the kernel of A . This means that all solutions of (*) have the form:

$$\hat{\beta} = X^\dagger y + z, \quad z \in \ker(X).$$

We claim that $\text{Im}(I_p - X^\dagger X) = \ker(X)$. First note that, $\text{Im}(I_p - X^\dagger X) \subseteq \ker(X)$, because

$$X(I_p - X^\dagger X)w = Xw - XX^\dagger Xw = Xw - Xw = 0.$$

And conversely, if $z \in \ker(X)$ then:

$$(I_p - X^\dagger X)z = z - X^\dagger Xz = z.$$

Thus $\ker(X) \subseteq \text{Im}(I_p - X^\dagger X)$. Hence

$$\hat{\beta} = X^\dagger y + (I_p - X^\dagger X)w,$$

as desired. \square

Theorem 5 (Minimum-norm Moore–Penrose solution). *Among all OLS estimator $\hat{\beta} \in \mathbb{R}^p$, the Moore–Penrose solution $X^\dagger y$ is the unique one of smallest l_2 -norm $\|\hat{\beta}\|_2$.*

Proof. This proof follows the proof of Corollary page 183 of [1].

By Theorem 4, any OLS minimiser can be written as

$$\hat{\beta} = X^\dagger y + (I_p - X^\dagger X)w.$$

We first show the two components are orthogonal:

$$\begin{aligned} (X^\dagger y)^\top (I_p - X^\dagger X)w &= (X^\dagger X X^\dagger y)^\top (I_p - X^\dagger X)w \\ &= (X^\dagger X X^\dagger y)^\top (I_p - X^\dagger X)w \\ &= (X^\dagger y)^\top (X^\dagger X)^\top (I_p - X^\dagger X)w \\ &= (X^\dagger y)^\top X^\dagger X (I_p - X^\dagger X)w \\ &= (X^\dagger y)^\top (X^\dagger X - X^\dagger X X^\dagger X)w \\ &= (X^\dagger y)^\top (X^\dagger X - X^\dagger X)w = 0. \end{aligned}$$

Thus they are orthogonal, and by Pythagoras:

$$\|\hat{\beta}\|_2^2 = \|X^\dagger y\|_2^2 + \|(I_p - X^\dagger X)w\|_2^2 \geq \|X^\dagger y\|_2^2.$$

The equality holds if and only if $(I_p - X^\dagger X)w = 0$, i.e. $w = 0$. In this case $\hat{\beta} = X^\dagger y$. Hence $X^\dagger y$ is the unique OLS estimator of smallest l_2 -norm. \square

Structure of the Minimum-Norm Estimator

To understand how to extend the Gauss–Markov perspective to the overparameterized setting, we begin by identifying which part of β is estimable when $p \geq n$ in the fixed-design linear model.

When $p \geq n$ and X has full row rank, the equation $\mathbb{E}[y] = X\beta$ no longer determines β uniquely. This non-uniqueness follows from a simple calculation: if $v \in \ker(X)$, then

$$X(\beta + v) = X\beta,$$

so every vector of the form $\beta' = \beta + v$ satisfies the same population equation. Because $\ker(X) \neq \{0\}$ when $p \geq n$, there are infinitely many such β' .

At this point it is natural to ask: which part of β can we hope to estimate? The earlier algebraic results in this paper already provide the answer. From our previous discussion of the Moore–Penrose pseudoinverse, we know that $X^\dagger X$ is the orthogonal projector onto $\text{rowsp}(X)$, and, among all vectors β' satisfying $X\beta' = \mathbb{E}[y]$, the one with smallest Euclidean norm is precisely $X^\dagger X\beta$.

Thus, rather than attempt to recover the full coefficient vector β , we focus on the component that is determined by the model, namely

$$\beta^* := X^\dagger X\beta.$$

This projected parameter is the unique minimum-norm solution to the population equation $X\beta' = \mathbb{E}[y]$ and represents exactly the part of β that influences predictions through X .

Moments of the Minimum-Norm Estimator

Before comparing the minimum-norm interpolator to other estimators in the overparameterized regime, we first examine its expectation and variance. These moments will be essential for later: the expectation identifies which parameter the estimator targets, and the variance formula will enter directly into the high-dimensional analogue of the Gauss–Markov theorem.

Expectation. A direct calculation yields

$$\mathbb{E}[\hat{\beta}] = X^\dagger \mathbb{E}[y] = X^\dagger X \beta = \beta^*.$$

Thus, in contrast to the classical regime (where unbiasedness means $\mathbb{E}[\hat{\beta}] = \beta$), the minimum-norm interpolator is unbiased for the projected parameter β^* . This observation will be crucial when defining the correct analogue of “unbiasedness” in the high-dimensional setting.

Variance. Since $\hat{\beta} = X^\dagger y$, we may write

$$\text{Var}(\hat{\beta}) = \text{Var}(X^\dagger y) = X^\dagger \text{Var}(y) (X^\dagger)^\top.$$

Under the fixed-design model $\text{Var}(y) = \Sigma$, so

$$\text{Var}(\hat{\beta}) = X^\dagger \Sigma (X^\dagger)^\top.$$

Together, these expressions show that $X^\dagger y$ behaves analogously to the classical OLS estimator, but with the full parameter β replaced by its identifiable component β^* . In the next subsection, we use these moments to characterize the appropriate notion of unbiasedness when $p \geq n$.

GAUSS–MARKOV ACROSS REGIMES

A central question in the theory of linear regression is that of optimality: among the class of valid estimators, which one minimizes variability? While the algebraic derivation of the minimum-norm solution $\hat{\beta} = X^\dagger y$ provides a geometrically intuitive candidate for the overparameterized regime, its statistical justification requires a more rigorous treatment. In this section, we extend the logic of the classical Gauss–Markov theorem to the high-dimensional setting. We begin by revisiting the traditional optimality results for $p < n$, and illustrating why they fail when $p > n$. We then construct a framework for unbiasedness that allows us to define the behaviour of $\hat{\beta}$ compared to other linear unbiased estimators of β^* .

Gauss–Markov in the Classical Regime

In the previous sections we have seen that, when $p \geq n$ and X has full row rank, the least-squares problem

admits infinitely many solutions and that the choice $\hat{\beta} = X^\dagger y$ singles out a particular interpolating estimator with a special geometric property: it has the smallest Euclidean norm among all minimizers of the OLS loss.

To better understand how unusual this overparameterized behaviour is, we now step back to the classical regime $p < n$. In this setting the design matrix X has full column rank, so the normal equations yield a unique solution. The OLS estimator is therefore uniquely defined, and as a consequence of full column rank, it also coincides with the minimum-norm solution.

In this setting we can meaningfully compare all linear and unbiased estimators of β and ask which one has the smallest possible variance. The Gauss–Markov theorem provides a complete answer. Under the homoscedastic linear model, the OLS estimator has the smallest covariance matrix among all linear unbiased estimators. This classical optimality result is stated below.

Theorem 6 (Gauss–Markov, classical regime). *Under the assumptions above and $p < n$, the OLS estimator $\hat{\beta}$ is the Best Linear Unbiased Estimator (BLUE) of β . More precisely, if $\tilde{\beta} = Ay$ is any other estimator which is linear in y and unbiased, (i.e. $\mathbb{E}[\tilde{\beta} | X] = \beta$), then*

$$\text{Var}(\tilde{\beta} | X) \succeq \text{Var}(\hat{\beta} | X).$$

Proof. This theorem proof can be seen from multiple sources including the following [4]. \square

This result relies critically on the full column rank of the design matrix X . Once we move to the setting $p \geq n$, this assumption fails, and several consequences follow immediately: β is no longer uniquely determined by the model, unbiased estimation of β is not well-defined, and the OLS estimator is no longer unique.

To understand what replaces the classical Gauss–Markov picture in the overparameterized regime, we begin by examining what the minimum-norm interpolator $X^\dagger y$ is actually estimating when $p \geq n$.

Notion of Unbiasedness for the Overparameterized Regime

In the classical regime, an estimator $\tilde{\beta}$ is called unbiased if $\mathbb{E}[\tilde{\beta}] = \beta$. When $p \geq n$, however, the parameter β itself is not determined by the equation $\mathbb{E}[y] = X\beta$, so this definition is no longer appropriate. The moments computed above suggest a more natural requirement: since β^* is the component of β that the model determines, and since $\hat{\beta} = X^\dagger y$ satisfies $\mathbb{E}[\hat{\beta}] = \beta^*$, it is reasonable to focus on estimators that are unbiased for β^* . We define this notion of unbiasedness as an estimator being row-space-unbiased for the parameter β^* .

That is, the estimator $\tilde{\beta}$ is rowspace-unbiased for β^* if $\mathbb{E}[\tilde{\beta}] = \beta^*$.

Another way to express unbiasedness is to demand that an estimator preserve the correct mean response, that is,

$$\mathbb{E}[X\tilde{\beta}] = X\beta.$$

This is the class of minimum-residual (interpolating) unbiased estimators of β^* .

The following result, adapted from Proposition 3 in [2], defines two classes of estimators. The first class is shown to be a proper subset of the second in the fixed-design model.

Theorem 7 (Unbiasedness for β^*). *Let $X \in \mathbb{R}^{n \times p}$ and let $\beta^* = X^\dagger X\beta$ be the projected parameter defined above. Consider an estimator $\tilde{\beta} \in \mathbb{R}^p$ under the fixed-design linear model. Then:*

(a) *If $\mathbb{E}[\tilde{\beta}] = \beta^*$, then*

$$\mathbb{E}[X\tilde{\beta}] = X\beta.$$

(b) *Conversely, if $\mathbb{E}[X\tilde{\beta}] = X\beta$, then*

$$X^\dagger X \mathbb{E}[\tilde{\beta}] = \beta^*.$$

Proof. Suppose (a), then we have

$$\begin{aligned} \mathbb{E}[\tilde{\beta}] &= \beta^* \\ X\mathbb{E}[\tilde{\beta}] &= XX^\dagger X\beta \\ \mathbb{E}[X\tilde{\beta}] &= X\beta. \end{aligned}$$

To prove (b), suppose that

$$\mathbb{E}[X\tilde{\beta}] = X\beta.$$

Then we have

$$\begin{aligned} X^\dagger X \mathbb{E}[\tilde{\beta}] &= X^\dagger \mathbb{E}[X\tilde{\beta}] \\ &= X^\dagger X\beta \\ &= \beta^* \end{aligned}$$

□

Note that for any linear estimator $\tilde{\beta} = Ay$ of β^* such that $\mathbb{E}[\tilde{\beta}] = \beta^*$, we have $A = X^\dagger$. To show this, we have

$$\begin{aligned} \mathbb{E}[\tilde{\beta}] &= \mathbb{E}[Ay] \\ &= AX\beta \\ &= X^\dagger X\beta \\ \implies AX &= X^\dagger X \\ \implies AXX^\dagger &= X^\dagger XX^\dagger \\ \implies A &= X^\dagger. \end{aligned}$$

Hence, the minimum-norm estimator $\hat{\beta}$ of β^* is the unique rowspace unbiased estimator when X is full row rank. If we were to apply Gauss-Markov to $\hat{\beta}$ and the class of rowspace unbiased estimators, it would be trivial and provide no discussion. So, we instead focus on comparing $\hat{\beta}$ to the larger class of interpolating unbiased estimators.

High Dimensional Gauss–Markov Theorem

We now return to the question of optimality. In the classical regime $p < n$, the Gauss-Markov theorem tells us that under a homoscedastic model, the OLS estimator is the linear unbiased estimator of β with smallest variance. When $p \geq n$, β is no longer unique. Furthermore, as shown following Theorem 7, the minimum-norm estimator $\hat{\beta}$ is the *unique* linear rowspace-unbiased estimator, so it is a well-defined problem to compare it to the broader class of interpolating estimators.

The next theorem is motivated by part (b) of Theorem 4 in [2], but is a stronger result. The result as stated in Theorem 8 is not in [2] or any other references to our knowledge.

By analyzing the variance over the geometrically distinct row space and null space of X , we discover that the minimum-norm estimator acts as the Best Linear Unbiased Estimator (BLUE) over these orthogonal subspaces of \mathbb{R}^p .

Theorem 8 (Gauss–Markov analogue in the overparameterized regime). *Assume the fixed-design linear model with homoscedastic noise $\text{Var}(\varepsilon) = \sigma^2 I_n$, and suppose that $X \in \mathbb{R}^{n \times p}$ has full row rank. Let $\hat{\beta} = X^\dagger y$ be the minimum-norm interpolator, and let $\tilde{\beta} = Ay$ be any linear estimator satisfying the interpolating condition $\mathbb{E}[X\tilde{\beta}] = X\beta$.*

Then:

(i) *For every vector $u \in \text{rowspace}(X)$,*

$$u^\top \text{Var}(\hat{\beta})u = u^\top \text{Var}(\tilde{\beta})u.$$

and for every vector $w \in \text{null}(X)$,

$$w^\top \text{Var}(\hat{\beta})w \leq w^\top \text{Var}(\tilde{\beta})w.$$

That is, $\hat{\beta}$ has the same variance over the rowspace of X and minimum variance over the nullspace of X when compared to all other interpolating estimators.

(ii) *Among all such estimators, $\hat{\beta}$ has the smallest total variance,*

$$\text{tr}(\text{Var}(\hat{\beta})) \leq \text{tr}(\text{Var}(\tilde{\beta})).$$

Proof. First, we note that for any deterministic matrix $M \in \mathbb{R}^{p \times n}$ we have

$$\text{Var}(My) = \sigma^2 MM^\top. \quad (1)$$

Now, we note that the set of right inverses of X is defined as

$$\mathcal{S}_R = \{A : A = X^\dagger + N, \text{ where } XN = 0\}.$$

Next, we define the set of linear interpolating estimators as

$$\mathcal{B} = \{\tilde{\beta} = Ay : A \in \mathcal{S}_R\},$$

since, for any $\tilde{\beta} \in \mathcal{B}$,

$$\begin{aligned} \mathbb{E}[X\tilde{\beta}] &= \mathbb{E}[XAY] \\ &= \mathbb{E}[XAX\beta + XA\varepsilon] \\ &= XAX\beta \\ &= X(X^\dagger X + NX)\beta \\ &= XX^\dagger X\beta + XNX\beta \\ &= X\beta. \end{aligned}$$

Now, using (1) given any $A \in \mathcal{S}_R$, we have

$$\begin{aligned} \text{Var}(Ay) &= \sigma^2 AA^\top \\ &= \sigma^2 (X^\dagger + N)(X^\dagger + N)^\top \\ &= \sigma^2 (X^\dagger N^\top + X^\dagger (X^\dagger)^\top + NN^\top + N(X^\dagger)^\top). \end{aligned}$$

Then, for any $\tilde{\beta} \in \mathcal{B}$, we have

$$\begin{aligned} \text{Var}(\tilde{\beta}) &= \text{Var}(Ay) \\ &= \sigma^2 (X^\dagger N^\top + X^\dagger (X^\dagger)^\top + NN^\top + N(X^\dagger)^\top) \\ &= \sigma^2 (X^\dagger (X^\dagger)^\top + NN^\top) + \sigma^2 (X^\dagger N^\top + N(X^\dagger)^\top) \\ &= \text{Var}(\hat{\beta}) + \sigma^2 NN^\top + \sigma^2 (X^\dagger N^\top + N(X^\dagger)^\top). \end{aligned}$$

Take any $u \in \text{rowspace}(X)$, and $w \in \text{null}(X)$, then we have

$$\begin{aligned} u^\top \text{Var}(\tilde{\beta})u &= u^\top \text{Var}(\hat{\beta})u + \sigma^2 u^\top NN^\top u \\ &\quad + \sigma^2 u^\top (X^\dagger N^\top + N(X^\dagger)^\top)u \\ &= u^\top \text{Var}(\hat{\beta})u, \end{aligned}$$

since $XN = 0 \implies N^\top u = 0$ for $u \in \text{rowspace}(X)$, as well as

$$\begin{aligned} w^\top \text{Var}(\tilde{\beta})w &= w^\top \text{Var}(\hat{\beta})w + \sigma^2 w^\top NN^\top w \\ &\quad + \sigma^2 w^\top (X^\dagger N^\top + N(X^\dagger)^\top)w \\ &= w^\top \text{Var}(\hat{\beta})w + \sigma^2 \|N^\top w\|^2 \\ &= \sigma^2 \|N^\top w\|^2, \end{aligned}$$

since $(X^\dagger)^\top w = 0$ for $w \in \text{null}(X)$. The second result is given by

$$\begin{aligned} \text{tr}(\text{Var}(\tilde{\beta})) &= \sigma^2 \text{tr}(X^\dagger N^\top + X^\dagger (X^\dagger)^\top + NN^\top + N(X^\dagger)^\top) \\ &= \sigma^2 (\text{tr}(X^\dagger N^\top) + \text{tr}(X^\dagger (X^\dagger)^\top) \\ &\quad + \text{tr}(NN^\top) + \text{tr}(N(X^\dagger)^\top)) \\ &= \sigma^2 (\text{tr}(N^\top X^\dagger) + \text{tr}(X^\dagger (X^\dagger)^\top) \\ &\quad + \text{tr}(NN^\top) + \text{tr}((X^\dagger)^\top N)) \\ &= \sigma^2 (\text{tr}(X^\dagger (X^\dagger)^\top) + \text{tr}(NN^\top)) \\ &= \text{tr}(\text{Var}(\hat{\beta})) + \sigma^2 \text{tr}(NN^\top) \\ &\geq \text{tr}(\text{Var}(\hat{\beta})), \end{aligned}$$

since $N^\top X^\dagger = (X^\dagger)^\top N = 0$. \square

This result updates the classical Gauss-Markov theorem by fundamentally splitting the feature space. Because we are comparing against the broader class of interpolating estimators ($XN = 0$), we can no longer claim minimum variance in every arbitrary direction v due to surviving cross-terms. However, Theorem 8 reveals that over the row space of X , every interpolating estimator performs identically. The optimality of the minimum-norm estimator is achieved entirely over the null space of X , where it maintains zero variance.

Corollary 9. For any vector $u \in \text{rowspace}(X)$, $u^\top (\text{Var}(\hat{\beta}) - \text{Var}(\tilde{\beta}))u = 0$, and for any vector $w \in \text{null}(X)$, $w^\top (\text{Var}(\hat{\beta}) - \text{Var}(\tilde{\beta}))w \leq 0$ for any alternative interpolating estimator $\tilde{\beta}$.

This emphasizes that there is no statistical advantage to be gained when predicting the identifiable signal. Any alternative interpolator introduces excess covariance matrix trace ($\sigma^2 \text{tr}(NN^\top)$) into the null space without improving signal recovery.

3 SIMULATIONS

This section goes through the simulations we ran to visualize and verify our findings related to the minimum-norm OLS estimator $\hat{\beta} = X^\dagger y$. Recall that our proofs require the fixed design model where the design matrix X is deterministic. Our simulations instead use a stochastic design matrix X and Monte-Carlo runs to determine the average behaviour of $\hat{\beta}$. All simulations that are done with Monte-Carlo use 50 iterations. We use a homoscedastic linear model

$$y = X\beta + \varepsilon, \quad \varepsilon \sim \mathcal{N}(0, 1),$$

where X and ε are randomly sampled for each iteration. In the simulations, we used $n = 50$, and $p = 200$, and we fix $\beta = 1_p$, the true linear coefficient. We will also be referring to the linear unbiased estimators $\tilde{\beta} \neq \hat{\beta}$ as *alternative estimators*.

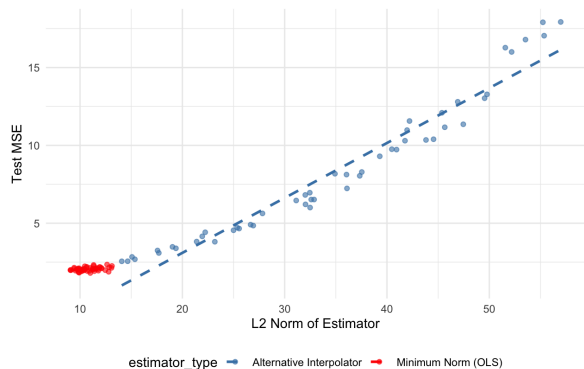


Figure 1: The test MSE of the estimators vs their norms.

3.1 Model Design

We used the a single design, which is simple and should highlight the behaviour of $\hat{\beta}$ nicely.

- *Normal model.* X is a random matrix whose entries are iid from a normal distribution with parameters $\mu = 0$, and $\sigma^2 = \frac{1}{p}$.

We scale the normal distribution by a factor of $\frac{1}{\sqrt{p}}$ to prevent the signal-noise ratio from becoming too large.

3.2 Simulation 1: Comparing the Test MSE between Estimators

Setup

The training data is generated by sampling X and ε , then building y . We use this X to create $\hat{\beta}$ as well as $\tilde{\beta}$. The alternative estimator $\tilde{\beta}$ is drawn from the set \mathcal{B} used in the proof of Theorem 8. That is, $\tilde{\beta} = Ay$, where $A = X^\dagger + N$, with $XN = 0$. This ensures the alternative estimator perfectly interpolates the data (minimum-residual), allowing us to visualize the variance leak into the null space. We then generate a test dataset of size 1000 to compute the test MSE, $\frac{1}{1000} \|y_{test} - \hat{m}(X_{test})\|^2$. This is done for each Monte-Carlo iteration.

Results

The MSE for the minimum-norm OLS estimator is less than that of the alternative estimator in each iteration. In figure 1, we can see that every alternative estimator has a larger norm than $\hat{\beta}$, and always has a larger test MSE. As well, the min norm solutions have very similar (and small) test MSE, but varying norm. The difference in norm can be attributed to the noise of X and ε , but note the predicting power of $\hat{\beta}$ is unaffected.

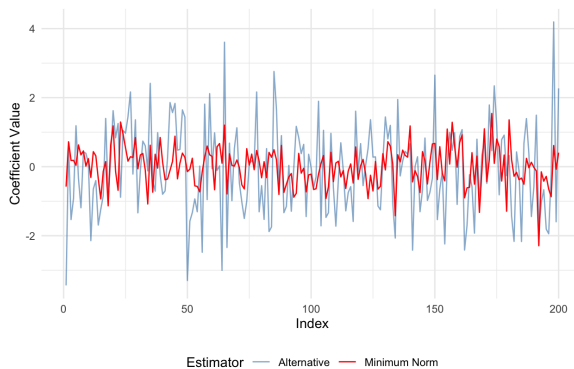


Figure 2: The noise of the coefficients of the estimators.

3.3 Simulation 2: Visualizing the Variance of the Estimators

Setup

We sample X and ε , and construct our dataset as normal. Then, we use this dataset find $\hat{\beta} = X^\dagger y$ and $\tilde{\beta} = My$. To visualize the noise, we look at $\hat{\varepsilon} = X^\dagger \varepsilon$ and $\tilde{\varepsilon} = M\varepsilon$. We only use a single instance of X and ε , with no Monte-Carlo iterations.

Results

In figure 2, the coefficients of $M\varepsilon$ show much more variance than that of $X^\dagger \varepsilon$. This aligns with our idea that the minimum-norm should be less noisy than the other unbiased estimators, since other unbiased estimators have non-zero variance over $null(X)$.

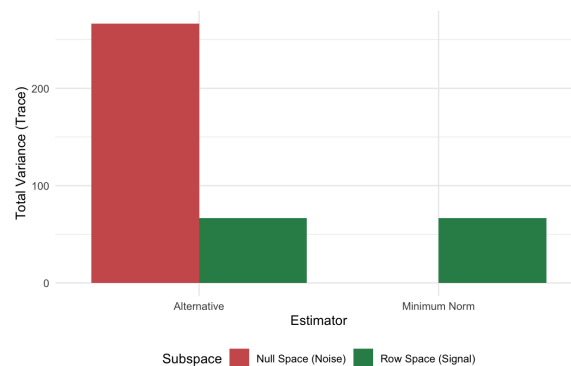


Figure 3: The total variance of the unbiased linear estimators of β^* .


PUZZLES

Here are some more puzzles by the editorial team, enjoy!

HEXADOKU

	11						16	13	6	5				3	
3						11			2	12	15				7
	16	15	5	8	6			11	4			2	13	14	1
	7			4		3	2					11		15	9
			7	12						6			2	8	3
	1				4	2		16		10				13	14
	2							5		8			7	4	16
	12	10				8			13		9	1		11	5
	13							9	3	15				12	11
		7						1				13	9	5	
12			16	1				7			5		4	2	8
4		11	10	15	16	9	12					7			6
9	15								14		8				12
5	10			6		13	14	15	9				1		
6			13	16	12	1	3	10				14		9	
7	14					5	9			4		6	8		13

MAZE



CREDITS

In alphabetical order

Editors-in-Chief

- Aahaan Rawal
- Hy Vu

Editors

- Aditya Sharma
- Louis Meunier
- Nitya Khirwar

Social Media Representative

- Callie Taylor

The Graduate Review Board

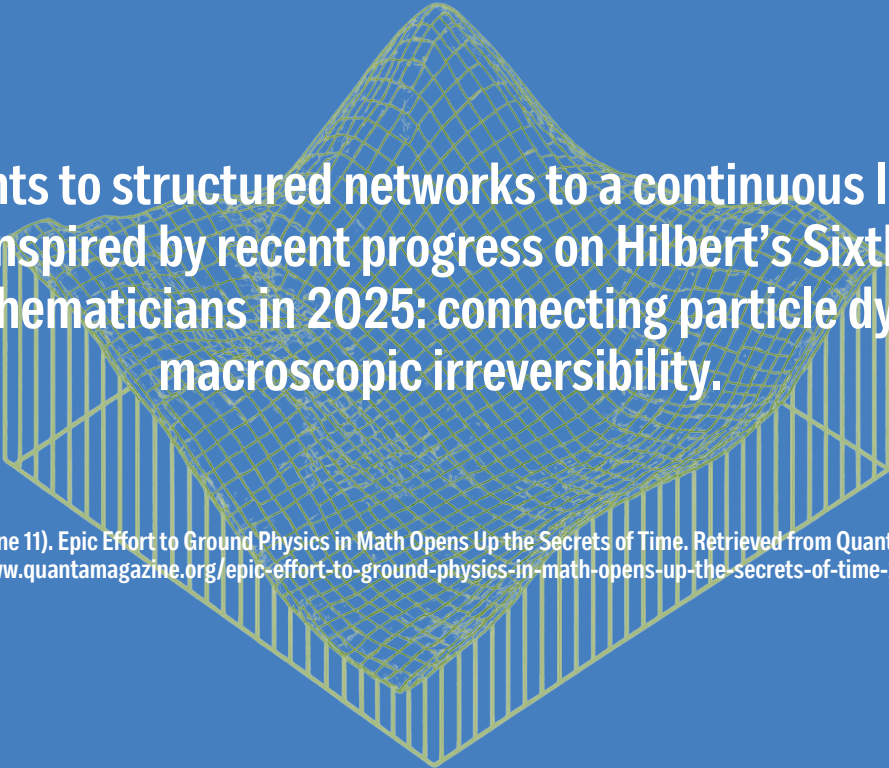
- Adrien da Silva
- William Holman-Bissegger
- Vincent Painchaud

Director

- Nicholas Hayek

ACKNOWLEDGEMENTS

We appreciate all the authors for sharing their hard work with the mathematical community and thank the entire graduate review board for volunteering their time to peer-review all articles for mathematical accuracy. Special thanks to Anush Tserunyan and Robert Robere for agreeing to share their wise words and advice in interviews with *The Delta Epsilon*. We are grateful to SUMS and AUS for their generous financial support to make the physical printing of this journal possible. And of course, we thank the readers of this journal, without whom this publication would be futile.



Scattered points to structured networks to a continuous landscape. The cover art is inspired by recent progress on Hilbert's Sixth Problem by three mathematicians in 2025: connecting particle dynamics to macroscopic irreversibility.

Sloman, L. (2025, June 11). Epic Effort to Ground Physics in Math Opens Up the Secrets of Time. Retrieved from Quanta Magazine website: <https://www.quantamagazine.org/epic-effort-to-ground-physics-in-math-opens-up-the-secrets-of-time-20250611/>