

Modeling Earth Systems Across Scales: Emergent Dynamics in Granular Materials  
and Inference-Driven Scales in Volcanic Records

by

Christopher Harper

A dissertation accepted and approved in partial fulfillment of the  
requirements for the degree of  
Doctor of Philosophy  
in Earth Sciences

Dissertation Committee:

Josef Dufek, Chair

Leif Karlstrom, Core Member

Thomas Giachetti, Core Member

Diego Melgar, Core Member

Chris Sinclair, Institutional Representative

University of Oregon

Summer 2025

© 2025 Christopher Harper  
All rights reserved.

This work, including text and images of this document, is licensed under a Creative  
Commons  
**Attribution-NonCommercial 4.0 International License.**



## DISSERTATION ABSTRACT

Christopher Harper

Doctor of Philosophy in Earth Sciences

Title: Modeling Earth Systems Across Scales: Emergent Dynamics in Granular Materials and Inference-Driven Scales in Volcanic Records

This dissertation explores emergent dynamics in granular materials and inference-driven scales in volcanic records, addressing the crucial role of scale in Earth sciences. It employs minimal models to reveal overlooked behaviors and assess the limits of data-driven inference.

Chapters 2 and 4 investigate granular materials, focusing on how particle deformation (softness) dictates system behavior. Contrary to traditional assumptions of infinite stiffness, this work shows that even minimal softness significantly alters force transmission via "compression chains" where particles deform in series. A deformable collision model was developed to quantify how these chains modify effective stiffness and reduce the impulse felt by a mobile grain. This directly explains thin layer stability in frictionless granular flows, proving compression chains as a necessary and sufficient mechanism. The model was generalized to heterogeneous systems, uniquely resolving multi-particle elastic collisions and revealing "contact ideals" in homogeneous systems. It also demonstrated that particle ordering and property variance profoundly impact momentum propagation in heterogeneous systems.

Chapter 3 shifts to the volcanic record, examining how data incompleteness and variable resolution affect inference. A statistical framework is introduced, employing temporal aggregation to extract reliable insights from sparse eruption records. This framework uses Freedman-Diaconis bin widths to identify statistically viable record segments, defining an "information resolution" for each. An "equivalent Poisson process" standardizes eruption intensity comparisons across varying resolutions, ensuring robust analysis. The methodology quantifies inferential power to determine if sufficient information exists to reject randomness, enabling rigorous comparisons of temporal structure. Lastly, it facilitates cross-volcano comparisons, even with vastly different record qualities, by normalizing for bin density, allowing for a "rhythm"

(dimensionless intensity) to characterize eruption "burstiness" or dispersal irrespective of observational regimes.

In essence, this dissertation provides novel physical insights into granular dynamics through explicit deformation modeling and a robust statistical framework for navigating the inherent limitations of Earth system records, thereby advancing our understanding of scale-dependent phenomena in both domains.

This dissertation contains both published and unpublished co-authored results.

## ACKNOWLEDGEMENTS

I would like to thank my advisor, Josef Dufek, for launching my graduate career and then giving me the freedom to pursue my own line of inquiry at my own (slow) pace. I have also been fortunate to receive outstanding mentorship throughout this process. Eric Breard and Laila Ennia have been exemplary mentors and role models, both in and out of the lab. Leif Karlstrom's style of inquiry sets an aspirational goal for my own. Finally, I am deeply grateful to my family and friends, who have, at times, acted as mentors and advisors—but most importantly, have been friends. Among these friendships, I want to especially thank my partner, Emily Doerner, without whom I'm not sure I would have ever finished.

To E.D.,  
my partner in everything.

## TABLE OF CONTENTS

Chapter	Page
I. INTRODUCTION . . . . .	17
II. THE ROLE OF COMPRESSIONAL DYNAMICS IN SETTING THE SCALE-DEPENDENT RHEOLOGY OF GRANULAR FLOWS: AN EXPLANATION FOR THE EMERGENCE OF THIN LAYER STABILITY . . . . .	19
2.1. Introduction . . . . .	19
2.2. Background . . . . .	21
2.2.1. Thin Layer Stability . . . . .	21
2.2.2. The Quartier Model: a simple theoretical framework . . . . .	22
2.3. Methods . . . . .	28
2.3.1. The Modified Quartier Model . . . . .	28
2.3.2. Modeling Deformable Particle Collisions . . . . .	31
2.4. Results . . . . .	35
2.5. Discussion . . . . .	40
2.6. Conclusion . . . . .	43
III. COMPLETENESS OF THE VOLCANIC RECORD . . . . .	45
3.1. Intro . . . . .	45
3.2. Background . . . . .	46
3.3. Methods . . . . .	48
3.3.1. Selecting Informative Sections of the Record . . . . .	49
3.3.2. Analyzing a portion: discretization and limits of information . . . . .	52
3.3.3. Quantifying Eruption Intensity via an Equivalent Poisson Process . . . . .	53
3.3.4. Comparing Structure and Inferential Power . . . . .	55

Chapter	Page
3.3.5. Comparing Records of Different Duration . . . . .	60
3.3.6. Summary of Methods . . . . .	60
3.4. Results . . . . .	61
3.4.1. R1: Selecting Statistically Viable Portions of the Record . . .	61
3.4.2. R2: Comparisons of intensity . . . . .	65
3.4.3. R3: Comparing Intensity Under Matched Bias (Temporal Resolution) . . . . .	67
3.4.4. R4: Comparing Records Under Matched Variance . . . . .	71
3.4.5. R5: Comparing records across drastically different qualities of recording. . . . .	73
3.5. Discussion . . . . .	76
3.6. Conclusion . . . . .	78
IV. DEFORMABLE COLLISION MODEL . . . . .	79
4.1. Introduction . . . . .	79
4.2. Theoretical Framework . . . . .	80
4.2.1. Generalizing the Equations of motion and resolving particle boundaries . . . . .	82
4.2.2. Identifying Chains of Compression . . . . .	84
4.3. Model Implementation and Software . . . . .	86
4.3.1. Overall Architecture and Workflow . . . . .	86
4.3.2. Core Algorithmic Steps: Implementing the Novelty . . . . .	87
4.3.2.1. Identifying Compression Chains ( <code>find_chains</code> ) . . .	87
4.3.2.2. Resolving Boundaries and Calculating Forces ( <code>solve_b0</code> , <code>satisfies_newton_third_law</code> , <code>calc_forces</code> ) . . .	89
4.3.3. Software Tooling and Accessibility . . . . .	90
4.4. Simple Simulation Walkthrough . . . . .	90

Chapter	Page
4.4.1. Setting Up the Simulation . . . . .	90
4.4.2. Running the Collision Model . . . . .	90
4.4.3. Accessing and Visualizing Results . . . . .	91
4.4.3.1. Key Plots . . . . .	91
4.4.3.2. Animation . . . . .	92
4.5. Validations and Demonstrative Examples . . . . .	93
4.5.1. Contact Ideals . . . . .	95
4.5.2. Variable Properties . . . . .	96
4.6. Discussion . . . . .	100
4.7. Conclusion . . . . .	101

APPENDIX: SUPPORTING INFORMATION FOR CHAPTER 2: THE ROLE OF COMPRESSIONAL DYNAMICS IN SETTING THE SCALE-DEPENDENT RHEOLOGY OF GRANULAR FLOWS: AN EXPLANATION FOR THE EMERGENCE OF THIN LAYER STABILITY . . . . .	102
A.1. Resolving Boundaries from Center of Mass . . . . .	102
A.1.1. Recovering boundaries within a compressed chain . . . . .	102
A.1.2. Identifying chains of compressed particles from center of mass positions . . . . .	105
A.2. Picking a time Scale for Impulse Calculations . . . . .	106
A.3. The role of the restitution coefficient . . . . .	108

## LIST OF FIGURES

Figure	Page
1. Perrin’s observations of thin layer stability manifesting in systems of frictionless grains [Perrin et al., 2021]. The open circles represent results measured from experimental runs and the black is best fit estimate. Critical angle denotes an angel of inclination for the system and $h/d$ is a dimensionless height. . . . .	22
2. A schematic of the Quartier system; a single gravity driven frictionless grain slides down a substrate of identical but fixed grains at an incline of $\theta$ . . . . .	23
3. (a) The force diagrams for the mobile grain instantaneously before and (b) after shock. The dashed illustrates the shifting particle contact. $V'$ represents the velocity preshock, and is shown in (b) to emphasize the acceleration forced by system geometry. . . . .	24
4. Velocity time series for a frictionless grain sliding over a rough inclined substrate as produced from a MFIX simulation of our frictionless Quartier system. . . . .	27
5. The substrate is represented by a variable number, $n$ , of deformable grains, allowing for changes in height. Offset A shows the one-dimensional deformable collision model that adjusts the contact force generated at the interface (illustrated by the greyed-out column), capturing how force propagation through sequential grain contacts modifies the overall force profile. . . . .	28
6. Velocity and deformation outputs from the deformable collision model. The left and right side depict two and three particle results respectively. . . . .	33
7. The percent error in deformations for a two particle system as a function of stiffness ( $k$ ). Percent error is calculated between the maximum deformation predicted from conservation equations and the maximum deformation observed in our model. The minimum stiffness used in our simulations is 100,000 N/m. . . . .	34

Figure	Page
8. The dashed lines trace out the magnitude of forces felt by a mobile grain, under the influence of gravity, as it collides into a vertical chain of particles of various length. Forces are normalized such that proportional comparisons between various chain lengths is accurate. The solid line indicates the number of additional particles currently deforming in series with the mobile grain (chain length). Time is non-dimensionalised such that the two particle interaction is one unit of dimensionless time. . . . .	37
9. An illustration of how our non-dimensionalization collapses results from simulations run with rigidities separated by an order of magnitude. . . . .	38
10. Blue, green, and red curves show impulse calculations for simulations run at rigidities of $10^4$ , $10^5$ , and $10^6$ respectively. Lighter and darker shades of each color represent initial velocities of 2 m/s and 4 m/s respectively. . . . .	39
11. Circles are our analytical results for critical angle as a function of non-dimensionalized height (eq 2.6). Triangles and dashed line are Perrin’s experimental results and line of best fit . . . . .	40
12. Visualization of contact network formed from two-dimensional polydisperse disks as shown in Papadopoulos et al. [2018]. . . . .	41
13. Scaling of critical angle produced by modifying the effective shock (eq2.10) according to the results of our energy budget analysis. The horizontal dashed line represents the 10-particle lengthscale cut off (a common threshold for mesoscale). . . . .	43
14. The x-axis shows the percentage of the record included, expressed as a fraction of total events. The left y-axis corresponds to the Freedman–Diaconis bin width ( $h$ ) of interarrival times, shown in blue. The right y-axis corresponds to the estimated event density, shown in purple. Vertical black lines indicate the locations where stable regions are detected. . . . .	51
15. Top panel: Histogram of interarrival times for events marked in the bottom panel. Bin width ( $\approx 0.1$ time units) is determined using the Freedman–Diaconis rule to discretize the record. Bottom panel: Discretized event timeline; active bins (containing $\geq 1$ event) are marked with yellow lines, while empty bins remain unshaded. . . . .	53

Figure	Page
16. fig:Left-hand-side: The top time series exhibits the highest Freedman Diaconis binwidth out of 100 simulations of $Uniform(\lambda = 10)$ over the unit interval. Middle represents the smallest binwidth observed from the same simulation set. The bottom record shows the higher resolution record being held to the same standard as the lower. Right-hand-side: From left to right, intensity and confidence interval estimates for top, middle, bottom . . . . .	56
17. Subplots iterate from reduced inference space of size 2, which results in 4 pairwise outcomes $\{(0, 0), (0, 1), (1, 0), (1, 1)\}$ to a reduced inference space of 10 which has $2^{10}$ possible outcomes. The left portion of each subplot shows the full distribution of possible outcomes on each reduced inference space, and the right subplot displays those outcomes distinguishable values as a CDF. The y-intercept on each subplot's CDF gives us the lowest possible significance level that can be achieved on the corresponding inference space. . . . .	59
18. fig: The forty-one Volcanoes in the GVP's Holocene eruption catalog with more than 50 events cataloged . . . . .	62
19. Each subplot contains the minimum FD binwidth plotted in blue and measured on the left axis, the density of events in purple measured on the right axis, all plotted against percent of record on the x axis. The dashed black lines indicate the end of "stable portions" and demarcate the portions of the volcanic record. . . . .	63
20. GVP only showing the records included up until the last stable portion ends for each volcano. . . . .	65
21. (a) The full record for each volcano is shown at high transparency, while the portion that comprises the first stable portion is shown opaque. (b) The same records but only showing the events included in the first stable portion and zoomed in to show a more representative time scale on the x-axis. . . .	66
22. (a) The full record for each volcano is shown at high transparency, while the portion that comprises the first stable portion is shown opaque. (b) The same records but only showing the events included in the first stable portion and zoomed in to show a more representative time scale on the x-axis. . . .	67

Figure	Page
23. (a) Intensity's calculated from the equivalent Poisson when $\Delta t$ is set to Izalco's resolution for each record (red) compared to the original optimal estimates (black). (b) Black lines represent a bin containing at least one eruption. All bins are aligned in time, such that the $i^{th}$ bin in each record corresponds to the exact same interval of time. . . . .	69
24. (a) (a) Intensity's calculated from the equivalent Poisson when $\Delta t$ is set to Izalco's resolution for each record (red) compared to the original optimal estimates (black). (b) Black lines represent a bin containing at least one eruption. All bins are aligned in time, such that the $i^{th}$ bin in each record corresponds to the exact same interval of time. . . . .	70
25. Histograms of inter-arrival times for 24 volcanoes under the maximum-variance condition. Orange subplots indicate insufficient inferential power to reject a randomness null hypothesis. Only records with adjusted bin widths $\geq$ their Freedman–Diaconis threshold are included. . . . .	72
26. Histograms of inter-arrival times under the median-variance condition, using Asosan as the benchmark. Only records with sufficient Freedman–Diaconis bin width are shown. Orange highlights indicate insufficient inferential power. . . . .	72
27. Dimensionless intensity comparison across volcanoes with differing record lengths and quality. Each point reflects a stable record segment normalized by bin density and mapped onto the unit interval. Color represents record duration (brighter = longer). . . . .	74
28. (a) Discretized records for Sheveluch and Kilauea at resolution which preserves bin-density. Active bins are yellow and the continuous process is shown as markers (b) The rhythm (dimensionless intensity) of Sheveluch and Kilauea. . . . .	75
29. Caption . . . . .	76
30. (a) Discretized records for Mazama and Villarrica at resolution which preserves bin-density. Active bins are yellow and the continuous process is shown as markers (b) Histograms of inter-arrival times with bin-width set to preserve bin-density. . . . .	76
31. Caption . . . . .	93

Figure	Page
32. (a) Deformation curves comparing a two-particle system (dashed lines) with a three-particle system, showing the change in system behavior with increased particle count. (b) Deformation curves for an 11-particle system, demonstrating continuity and resolution across multiple particles. (c) Zoomed-in view of an 11-particle system under gravity, illustrating the fine-grained continuity of deformation curves. (d) Full-time view of an 11-particle system under gravity, revealing a larger-scale oscillatory behavior characteristic of the system’s collective response. . . . .	94
33. These subplots illustrate the contact structure for particle one, showing the length of the compression chain in contact with particle one through time. These systems are unconstrained (last particle not fixed in space) and not subject to gravity. Time is non-dimensionalized to the timescale of contact for a two-particle system of equivalent properties. (a) The 2-particle system sets the non-dimensionalization of time. (b) In the 3-particle system, at most two other particles can be in contact with particle one. (c) 4-particle system. (d) For systems with 5 or more particles, this represents the limiting contact structure from the first particle’s point of view, with no further changes observed as the number of particles continues to increase. . . . .	98
34. Illustration of the complex contact structure behavior that arises in the presence of gravity. The y-axis represents contact length, and the x-axis represents non-dimensionalized time. . . . .	98
35. Impact of particle ordering on the final particle’s velocity. (a) Compares alternating and well-sorted stiffness distributions. (b) Compares alternating and well-sorted mass distributions. . . . .	99
36. Impact of material property variance on the final particle’s velocity. (a) Compares low and high variance stiffness distributions. (b) Compares low and high variance mass distributions. . . . .	99
A.37.A clear separation of mobile grain-particle and mobile grain-chain interactions develops as chain length increases. It arises similar to a separation of constructive interference. . . . .	108
A.38.Critical angles for simulations of two to seven particles when using contact structures to pick temporal bounds for impulse calculation. . . . .	109

Figure	Page
A.39. Visualization of contact structure (blue), forces felt (orange), and compression of the mobile grain for a 4 particle chain. The purple dashed line indicates where the particle is maximally compressed, which sets the timescale for our impulse calculation. . . . .	110
A.40. Sensitivity of our model to different values for the restitution coefficient .	110

## LIST OF TABLES

Table	Page
1. Parameters which arise from the Quartier models simplified geometry. . .	24
2. All bitmask combinations and their resulting partitions for $N = 4$ particles. . . . .	88

## CHAPTER I INTRODUCTION

The Earth is full of complex, beautiful systems, many of them playing out invisibly beneath our feet, others unfolding in dramatic, surface-level events. Within the Earth sciences, we try to make sense of these processes. This is no small undertaking. Not only are the systems themselves intricate and wide-ranging, but so too are the methods and tools we use to study them. Physics, chemistry, math, computer science, sometimes all at once. With so many ways to frame an Earth science question, it's easy to get swept up in complexity. This dissertation takes a different tack. Rather than chasing complete models of entire systems, it focuses on something smaller: digging into a few key processes and asking what can be learned from minimal assumptions. By building simple, physically grounded models to reveal overlooked behaviors, and assessing what kinds of inferences are even possible given the limits of data, this dissertation demonstrates the use of careful minimal models to answer some of the most illusive and persistent question in the Earth sciences, questions of scale.

Questions of scale merge in all scientific inquiry, but play a particularly central role in Earth science, where systems span vast ranges of space and time, and where direct observation or experimentation at scale is often impossible. The same process may appear entirely different depending on how closely we zoom in, or how long we observe it. In some cases, these differences reflect real physical transitions: new behaviors emerge, dominant forces shift, and previously valid assumptions begin to break down. These are questions of physical scale, classically illustrated in fluid dynamics through the transition to turbulence [Reynolds, 1883], and more recently in granular materials, where systems have been proposed to exhibit self-organized criticality [Bak et al., 1988]. In other cases, the change is not in the system but in our ability to observe it: sparse, uneven, or noisy data alter how the system appears—leading us to draw incomplete or even misleading conclusions. These are questions of inference scale. Sadler [1981]'s work in stratigraphy remains a touchstone here, demonstrating how observational gaps shape our estimates of depositional rates and constrain what we can meaningfully infer about the processes themselves.

In many Earth science contexts, these two types of scale challenges appear together. Seismology, for instance, must grapple with the completeness of seismic records [Wiemer and Wyss, 2000], a problem of inference, while simultaneously

exploring whether earthquakes follow a scale-free, critical behavior [Saichev and Sornette, 2006]—a question of physical scale. Across scientific history, examples like these have shown that the reliability of models often hinges on scale: understanding where our theories apply, and where they don't, is key to building robust, general insights. This is especially true in the Earth sciences, where natural systems resist direct control and simple extrapolations can lead us astray.

This dissertation takes up one example of each type of scale problem. Chapters two and four focus on granular materials, exploring how a simple physical mechanism, deformation, can lead to emergent behaviors that traditional models overlook. These chapters show how even small changes in model assumptions can produce large-scale consequences. Chapter three shifts to the volcanic record, asking how uncertainty, incompleteness, and variable resolution affect our ability to make meaningful inferences about volcanic processes. Together, these cases explore how scale, both physical and observational, shapes the boundaries of what we can know, and how we come to know it.

## CHAPTER II

### THE ROLE OF COMPRESSIONAL DYNAMICS IN SETTING THE SCALE-DEPENDENT RHEOLOGY OF GRANULAR FLOWS: AN EXPLANATION FOR THE EMERGENCE OF THIN LAYER STABILITY

From *C. Harper, J. Dufek, and E. Breard. The role of compressional dynamics in setting the scale-dependent rheology of granular flows: An explanation for the emergence of thin layer stability. Physical Review E. 111(6), 2025.*

#### 2.1 Introduction

Granular processes are as ubiquitous as they are complex, present in the bases of fast flowing pyroclastic density currents, and the slow shifting stability (creeping) of slopes. The governing dynamics of these fast and slow extremes have disparate characterizations; the former better described by kinetic theory and binary particle collision, while the latter is dominated by enduring grain contacts, friction, and slow deformation. Despite different characterizations, both dynamics can manifest simultaneously in even the simplest granular systems, as demonstrated in Rocha et al. [2019] experiments, in which flows of mostly mono-disperse inclined granular systems exhibit a fast flowing center guided by slowly shifting to nearly stationary edges. One of the great challenges of describing granular systems lies in finding the physical mechanisms capable of producing such complex rheologies, which, while essential for accurate modeling, are poorly understood. Much of this complexity stems from a scale-dependent rheology: even modest changes in system scale can fundamentally alter its behavior, making the appropriate analogy to solid, liquid, or gas phases of matter difficult [Jaeger et al., 1996]. This complexity is in part because the grains form a fabric of contacts which allows for grains, not directly in contact with one another, to significantly affect each other [Papadopoulos et al., 2018]. This lengthscale at which non-local particles can affect one another is referred to as the mesoscale, and the effects which manifest from particle contacts at the mesoscale are typically referred to as non-local effects.

Despite these complexities, an insightful argument by Da Cruz et al. [2005] showed that evolution of stresses in systems of simply sheared infinitely rigid grains are well characterized by a single dimensionless parameter, the inertial number " $I$ ". Surprisingly this global description for simple shear could be used to describe the evolution of stresses locally in inhomogeneous systems giving rise to a powerful

rheological frame work known as the local  $\mu(I)$ -rheology [Forterre and Pouliquen, 2008]. This framework has been successfully expanded to capture a wide range of dense granular systems [Boyer et al., 2011, Amarsid et al., 2017], however these approaches universally struggle with the scale-dependent dynamics which emerge in granular domains with low inertial numbers: such as transition to flow, jamming, creep, and many other non-local effects.

Phenomenological descriptions, such as the kinetic elasto-plastic (KEP) theorem [Bocquet et al., 2009] or granular fluidity [Henann and Kamrin, 2014], have been proposed to modify the  $\mu(I)$ -rheologic framework and capture quasi-static granular behavior. However, a kinematic state variable which can describe the physical underpinnings of these phenomenology remains elusive. In this work, we argue that the foundational assumption of infinite particle stiffness, central to the  $\mu(I)$ -rheology, must be relaxed to reveal these mechanisms. We posit that even an infinitesimal degree of softness, and hence deformation, permits “rigid” particles to compress in series, thereby introducing an effective stiffness that modifies the local effective friction coefficients of the end-members.

By “deformation” we refer to the elastic compression of a particle itself. We will show that even if this compression amounts to only one-millionth of a percent of the particle diameter, it can dramatically alter the transmission of contact forces between particles. For example, consider dropping a textbook onto two different spring systems, one consisting of a single spring and the other composed of two identical springs arranged in series. Although the springs are identical, the force transmission differs markedly between the two configurations. Regardless of how stiff the springs are, no matter how minuscule the absolute deformation, the ability of springs to compress in series persists, thereby reducing the overall effective stiffness. For springs of equivalent stiffness, this behavior is captured by the left hand side of proposition 2.1.

$$k' = \frac{k}{n} \longrightarrow \frac{|k - k'|}{k} = \frac{n - 1}{n} \quad (2.1)$$

The implication (right hand side) is that the percent change in effective stiffness depends solely on the number of springs in series and is independent of the absolute stiffness. In the limit of infinite stiffness, the effective stiffness equals that of an individual spring—rendering the two systems indistinguishable and erasing any observable difference due to serial compression. **We contend that a**

**similar phenomenon occurs in granular materials: even an infinitesimal deformation permits particles to compress in series, significantly altering local friction coefficients and stress transmission.** Capturing this effect in a manner that accurately reflects the nature of granular systems—where particles behave as separable, massive springs—is challenging, and addressing this challenge is the focus of our work

To validate the importance of serial compression (also referred to as compression chains) in producing previously unaccounted-for granular phenomena, we demonstrate that incorporating this effect into a simple kinematic model of inclined granular flow explains the emergence of a widely observed yet still incompletely understood phenomenon—thin layer stability. Thin layer stability refers to the observation that the critical stresses for the onset and cessation of flow depend on a non-dimensionalized system height (hereafter “height”), and it exhibits both scale dependence and non-locality. The non-dimensionalized height is derived from a system height normalized by average particle diameter, thus setting the scale-dependence to the “number of particles tall”. A thorough understanding of thin layer stability is essential for explaining many other perplexing granular phenomena; for instance, it plays a role in silo clogging [Kamrin et al., 2024] and in levee formation, which can lead to self channelization and increased run out observed in many large-scale and complex geophysical flows (e.g. lava flows, pyroclastic density currents, debris flows etc)[Mangeney et al., 2007, Sparks et al., 1976, Kokelaar et al., 2014, Jessop et al., 2012, Johnson et al., 2012].

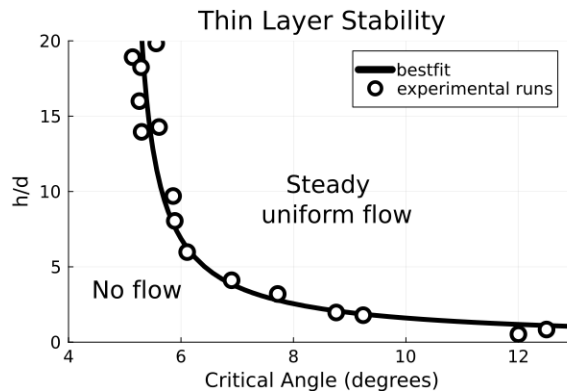
We capture the effects of deformation on the kinematics of inclined granular flow within a modified "grain on a sandpile model" [Quartier et al., 2000] coupled with our own deformable collision model. The results are compared to Perrin’s experimental results for inclined planar granular flow of frictionless particles [Perrin et al., 2021], and show an excellent recovery of the observed thin layer stability. We conclude with a discussion for how this phenomenon manifests more generally, and when coupled with knowledge of contact networks could provide a physical basis for the phenomenological notions of granular fluidity and the transport of granular temperature.

## **2.2 Background**

### **2.2.1 Thin Layer Stability**

Thin layer stability, an incompletely explained non-local effect, is most simply exhibited in inclined planar flow, where it manifests as a dependence between the

angle of repose and the non-dimensionalized system height [Pouliquen, 1999]. Thinner systems of grains have a higher angle of repose, with this critical angle decreasing asymptotically as system height increases. This phenomenon has a long history of experimental observations in a variety of granular systems, ranging from glass beads to walnut shells [Pouliquen and Renault, 1996]. This breadth of systems implies a general importance to mesoscale dynamics and significance to non-local effects in general, while the simplicity of systems in which it can manifest emphasize a link to fundamental granular behavior; thin layer stability has been observed in inclined planar flows of mostly mono-disperse frictionless grains (Figure 1).



*Figure 1.* Perrin’s observations of thin layer stability manifesting in systems of frictionless grains [Perrin et al., 2021]. The open circles represent results measured from experimental runs and the black is best fit estimate. Critical angle denotes an angel of inclination for the system and  $h/d$  is a dimensionless height.


**Thus the roots of this phenomenon should be captured in models which emphasize and exemplify the role particle geometry plays in setting system behavior, and as we show in the next two sections – they are.** This is done through modifying a simple yet powerful theoretical framework pioneered by Quartier et al. [2000], with our final results being benchmarked against Perrin et al. [2021] experimental finding (Figure 1).

### 2.2.2 The Quartier Model: a simple theoretical framework

The Quartier model consists of one mobile grain and a substrate of fixed grains on an inclined plane, as shown in Figure 2. For clarity, we refer to the mobile particle as a grain, and the fixed substrate as particles. A complete analysis of this system was done in Quartier’s publication on the “Dynamics of a grain on a sandpile”. Despite its simplistic nature the analysis clearly reveals promising roots for many complex

granular phenomenon, such as the ability for a substrate of grains to form periodic potential traps, how these traps, in turn, bifurcate the system into a static and dynamic steady state, and how simple geometric constraints can give rise to hysteresis (we will highlight this behavior later [Quartier et al., 2000]. The analysis performed here involves frictionless particles to be consistent with the experiments of Perrin et al. [2021], and can take a more simplified approach. Thus our governing equations (and the derivation method) are distinct from Quartier’s original work.

In this frictionless analysis, the grain does not roll, but slides over the surface of the glued particles as it is driven by gravity. It is imbued with an initial kinetic energy  $E_0$ , such that it never becomes airborne, and in this way the dynamics of the system are determined by the interactions at points A,B, and C.



*Figure 2.* A schematic of the Quartier system; a single gravity driven frictionless grain slides down a substrate of identical but fixed grains at an incline of  $\theta$ .

At point A the grain is sliding down the crest of one particle and into the gravity well at point B. The grain will then crest the next particle at point C. Whether the grain successfully crests the next particle depends on a balance between the kinetic energy gained from cresting the previous particle, the cost of escaping the energy well (C), and the energy lost in the shock experienced at point B ( $\mathcal{E}$ ). The first two are solely dependent on system geometry and summarized in table 1, the ability for the particle to remain mobile is determined by the interaction at point B, which bears further scrutiny.

Figure 3 shows the force diagram for the mobile grain right before and after the shock experienced at the bottom. During the shock, the velocity accelerates from being perpendicular to particle contact AC (panel(a)) to perpendicular to particle

Key system parameters	
Initial kinetic energy	$E_0$
Gravity well	$\mathcal{C} = mgr (1 - \cos(\frac{\pi}{6} - \theta))$
Kinetic energy gained per crest	$\Delta E = mgr \cdot \sin \theta$
Energy after first shock	$E_1 = (E_0 + \Delta E)\mathcal{E}$

Table 1. Parameters which arise from the Quartier models simplified geometry.

contact BC (panel (b)), and some energy is dissipated and captured by the term  $\mathcal{E}$  (kinetic energy after the shock/kinetic energy before the shock).

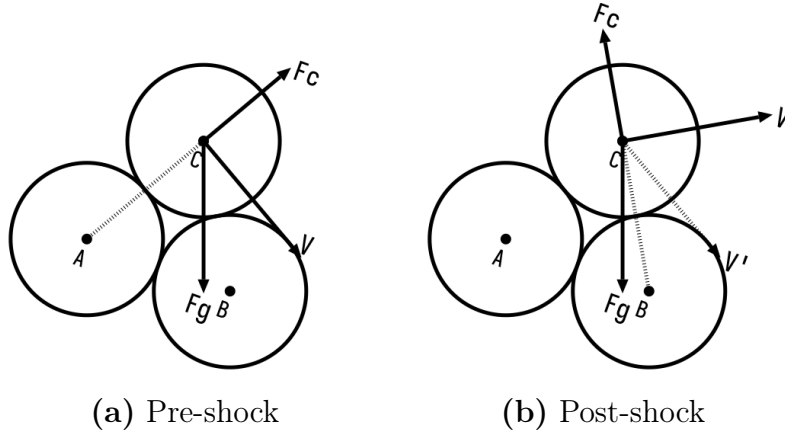


Figure 3. (a) The force diagrams for the mobile grain instantaneously before and (b) after shock. The dashed illustrates the shifting particle contact.  $V'$  represents the velocity preshock, and is shown in (b) to emphasize the acceleration forced by system geometry.

This allows for a powerful analysis of the system by expressing the energy at each successive shock in a recursive manner (last line of table 1), enabling the identification of steady states. Such steady states are crucial for understanding system stability. By analyzing these states, we can determine whether the system returns to equilibrium or continues to move, providing a key method for evaluating stability. This is accomplished as follows:

If the energy after the first shock is sufficient to escape another well ( $\mathcal{C} \leq E_1$ ), then we can describe the energy after the second shock as  $E_2 = (E_1 + \Delta E)\mathcal{E}$ . Expanding and simplifying for successive terms reveals the following pattern:

$$E_n = E_0 \cdot \mathcal{E}^n + \Delta E \cdot \mathcal{E} \underbrace{(\mathcal{E}^{n-1} + \mathcal{E}^{n-2} + \dots + \mathcal{E}^2 + \mathcal{E} + 1)}_{S_n} \quad (2.2)$$

The part labeled  $S_n$  can be simplified using the following series identity.

$$S_n = 1 + x + x^2 + \dots + x^n = \frac{1 - x^{n+1}}{1 - x} \quad (2.3)$$

Resulting in a useful closed form expression for the energy after the  $n^{th}$  shock:

$$E_n = E_1 \cdot \mathcal{E}^{n-1} + \Delta E \cdot \mathcal{E} \left( \frac{1 - \mathcal{E}^{n-1}}{1 - \mathcal{E}} \right) \quad (2.4)$$

If the energy gained from cresting a particle is greater than the energy dissipated through the shock ( $\mathcal{E} > 1$ ),  $E_n$  will go to infinity, and this grain will accelerate indefinitely. However, if the energy dissipated from the shock is greater than the energy gained, then the system will asymptotically approach a constant energy state:

$$\lim_{n \rightarrow \infty} E_n = \Delta E \left( \frac{\mathcal{E}}{1 - \mathcal{E}} \right) \quad (2.5)$$

We can set this limiting energy equal to the expression for the cost of the energy well to find a critical point of our system, the angle at which our grain can first enter into steady state motion. This minimum angle is called the critical angle:

$$\begin{aligned} \Delta E \left( \frac{\mathcal{E}}{1 - \mathcal{E}} \right) &\geq mgr \left( 1 - \cos \left( \frac{\pi}{6} - \theta \right) \right) \rightarrow \\ \frac{\mathcal{E}}{1 - \mathcal{E}} &\geq \frac{2 \sin^2 \left( \frac{\pi}{12} - \frac{\theta}{2} \right)}{\sin(\theta)} \end{aligned} \quad (2.6)$$

From equation 2.6 we see that the critical angle is only dependent on the angle of inclination and the dissipative term,  $\mathcal{E}$ , which represents the ratio of kinetic energy after and before each shock (equation 2.7).

$$\mathcal{E} = \left| \frac{v'_i - v_i}{v'_i} \right|^2 \quad (2.7)$$

The velocity instantaneously before and after the  $i^{th}$  shock is represented as  $v'_i$  and  $v_i$  respectively. Using the geometric constraints of our system, and introducing a new material property, the normal restitution coefficient ( $e_n$ ), we can solve for the velocity post-shock (equation 2.8).

$$v_i = \sqrt{\left(\underbrace{v'_i \cdot \frac{1}{2}}_{\hat{i}}\right)^2 + \left(e_n \left(\underbrace{v'_i \cdot \frac{\sqrt{3}}{2}}_{\hat{j}}\right)\right)^2} \quad (2.8)$$

This can be interpreted as the original velocity multiplied with the new component vector post shock, with  $\hat{i}$  being the direction tangent to new particle contact, and  $\hat{j}$  the direction perpendicular. In this form we can clearly see the role of the normal restitution ( $e_n$ ), which represents the percentage of velocity normal to particle contact conserved after collision. The factors of  $\frac{1}{2}$  and  $\frac{\sqrt{3}}{2}$  stem from the equilateral triangle created from the geometry of mono-disperse particles in a configuration shown in fig 3. By substituting this expression into  $v_i$  in equation 2.8 and canceling out the  $v'_i$  term we arrive at our full expression for shock (equation 2.9):

$$\mathcal{E} = \left|\frac{v_0 - v_s}{v_0}\right|^2 = \left(1 - \sqrt{\left(\frac{1}{2}\right)^2 + \left(e_n \frac{\sqrt{3}}{2}\right)^2}\right)^2 \quad (2.9)$$

To validate the power of the geometric analysis we turn to Discrete Element Modeling (DEM), where we can set up a system mirroring Figure 2 which matches all of our assumptions (monodisperse, frictionless, and constant contact). Because the particles are currently assumed to be perfectly rigid (no real deformation occurs) we have to set the restitution coefficient ( $e_n$ ) to zero to ensure the particles maintain constant contact with the substrate. This assumption of constant contact is also one made by Quartier et al. [2000] in his original works (this assumption will be relaxed in the subsequent section). In setting  $e_n$  to zero in equation 2.8 we get a value of  $\frac{1}{4}$ , and root finding equation 2.6 gives us the critical angle which bifurcates our system into either perpetual motion or rest is 10.21 degrees.

To perform our DEM validation we implemented the U.S. Department of Energy's MFIX-DEM approach for frictionless mono-disperse grains [Garg et al., 2012]. A two-dimensional rectangular domain is created with cyclic boundary conditions for the vertical edges of the domain. One layer of grains is fixed to the base by resetting the acceleration to zero at the end of each computation cycle. A single grain is placed in the well between the first and second particle and is set to have an initial velocity sufficient to escape the gravity well. The material parameters can be set such that the friction coefficient is zero and the normal restitution coefficient

( $e_n$ ) is nearly zero ( $10^{-20}$ ). Finally, the vector for gravity is inclined from vertical to reproduce the angle of inclination.

The minimum angle at which the grain can maintain steady state matches the critical angle calculated from finding the roots to equation 2.6 to one-hundredth of a degree, as shown in Figure 4.




Figure 4. Velocity time series for a frictionless grain sliding over a rough inclined substrate as produced from a MFiX simulation of our frictionless Quartier system.

While many elements of this model are idealized, particularly the demand that  $e_n = 0$ , we see that even in this simple form this model can explain a potential root cause for hysteresis. Namely, that the angle to maintain steady motion ( $\theta = 10.21$ ) is less than the angle needed to start moving from rest  $\theta > 30$ . This rest angle stems from the fact that at an inclination of thirty degrees, the mobile grain is balanced right at the top of particle B (Figure 3), and at any perturbation would begin to slide on its own.

**We conclude this section with a simple question. If this analytical framework exemplifies the role that particle geometry plays in setting system wide behavior, where in this framework does thin layer stability live?**


It lives in term 4 in table 1, the term which relates the shock to energy loss. Rather than having the shock be instantaneous, we will allow for deformation. The normal force between the substrate and the mobile grain (Figure 3) then acts as the geometric friction not just in some instantaneous shock, but as the particle crests out

of the well. However, since deformation is now allowed, this normal force (geometric fraction) then depends on a complex series of compressions as the impulse propagates through the substrate, which will in turn depend on how many particles deep is the substrate. To capture this effect we move to a further modified Quartier Model.

## 2.3 Methods

### 2.3.1 The Modified Quartier Model

Our modified Quartier model introduces two key modifications to the traditional setup: (1) the particles are allowed to deform and (2) we introduce a new independent variable  $n$ , which denotes, the number of layers of stacked particles that remain in contact above the substrate,  $n$ .



*Figure 5.* The substrate is represented by a variable number,  $n$ , of deformable grains, allowing for changes in height. Offset A shows the one-dimensional deformable collision model that adjusts the contact force generated at the interface (illustrated by the greyed-out column), capturing how force propagation through sequential grain contacts modifies the overall force profile.

As noted in the introduction, permitting deformation fundamentally alters the collisional forces between a mobile grain and the substrate. Recall our introductory example—dropping a textbook onto a series of springs. In a more technical and representative scenario, a mobile grain experiences a shock at the base of a gravitational well. Here, the contact force (shown in Figure 3) is governed by a length scale determined by the compressible bed, where the series of substrate particles (represented by the greyed-out particles in Figure 5) deform in series.

To capture this effect, we model a simple system of massive, deformable particles—analogueous to springs—as illustrated in offset A. The details of this model are presented in the next subsection; for now, we focus on how this new behavior alters the key geometric constraints outlined in the previous section.

Ultimately, allowing deformation alters particle contacts, thereby modifying the contact forces between the mobile grain and the substrate. This modification, in turn, affects the total impulse delivered while the grain is within the particle well. Once we determine how the impulse varies as a function of substrate height,  $n$  (as derived in the subsequent modeling section), we can compute a revised post-shock velocity and shock term, ultimately yielding a new dependence of the critical angle for steady state flow:

$$\begin{aligned}\alpha(n) &= \frac{J(n)}{J_0} \\ \mathcal{E}(n) &= \left| \frac{v_0 - v_s \cdot \alpha(n)}{v_0} \right|^2 \\ \frac{\mathcal{E}(n)}{1 - \mathcal{E}(n)} &\geq \frac{2\sin^2(\frac{\pi}{12} - \frac{\theta}{4})}{\sin(\theta)}\end{aligned}\tag{2.10}$$

Arriving at the final equation in equation set 2.10 requires the composition of two models. We begin with a straightforward model of deformable grains under gravity, as depicted in Figure 5(offset A), which captures the essential effect of serial compression discussed in the introduction. This model characterizes the impulse experienced by the first particle as a function of the number of particles,  $n$ , involved in the compression chain (three in the figure, for example). The resulting modified impulse (first equation in equation set 2.10) then informs the altered shock response in a correspondingly modified Quartier model with substrate thickness  $n$ , as illustrated by the interaction with the grayed-out gravity well and encapsulated in the second equation in equation set 2.10. With this modified shock incorporated into our Quartier framework, we can recover the global behavior described earlier and encapsulated in the final line of equation set 2.10. In turn, this equation delineates the dependence of the system’s critical angle—which governs system stability (i.e., whether steady state motion is possible)—thereby introducing a system-wide stability dependence on  $n$ . Our results will show that this scaling mirrors the thin layer stability observed in Perrin’s model.

This composite model incorporates several simplifying assumptions: (1) compression is modeled as a one-dimensional process, and (2) the substrate is assumed to have a uniform surface geometry. As we show in the next section, modeling deformable, massive springs presents significant challenges. Nonetheless, the primary objective of this study is to establish both the existence and the significance of this heretofore unaccounted-for phenomenon. Moreover, our first-order model not only suffices to demonstrate its existence but also yields results (see Figure 11) that compare favorably with observations from systems composed of monodisperse, yet more naturally packed, grains (Perrin et al. 2020)(RefPerrin2021).

Before proceeding to the next section, where we detail our methods for modeling deformable particles in series, we first present some important nondimensionalization and preliminary timescale analyses:

In many discrete element models, the normal contact force between two particles (denoted  $F_c$  in Figure 3) is determined by the material property known as the “normal stiffness coefficient,” denoted  $k_n$  to distinguish it from the material stiffness under shear. The shear stiffness is not relevant to our model set-up, so we drop the subscript and simply refer to this normal stiffness as  $k$ . A nondimensional stiffness, referred to as the rigidity ( $\kappa$ ), is defined by normalizing the stiffness with the confining pressure (eq 2.11):

$$\kappa = \frac{k}{P} \tag{2.11}$$

da Cruz (2005) notes that for  $\kappa \geq 10,000$ , there is little variation in coordination number with further increases to rigidity and the length scales for particle deformation become negligible to gaps between neighboring grains. Providing an example for physical intuition, a bead of glass at the bottom of a 50 cm deep pile, will have a rigidity of 40,000 [Da Cruz et al., 2005]. Our analysis will include rigidity values ranging three order of magnitude [10,000, 1,000,000], because our confining pressure for the mobile grain is determined by the gravitational body force ( $9.81 \text{ m/s}^2$ ) this maps to stiffnesses of 100,000  $N/m$  and 10,000,000  $N/m$ , respectively. The lower bound of stiffness is used for the timescale calculations as this will give us an upper bound for the timescales under consideration.

The timescale, which will dominate the Quartier model, is the time it takes for a grain to rise from particle well to crest (point b to c in Figure 2). This is estimated

by taking the minimum change in velocity needed to escape the well and dividing by the acceleration provided from the gravitational field (equation 2.12).

$$t_c = \frac{\Delta V}{a} = \frac{\frac{1}{4}\sqrt{2g(2r(1 - \sin(60)))}\cos(60)}{g} \approx 0.41s \quad (2.12)$$

We use the equation for rigid particle contacts to estimate the timescale for deformation [Garg et al., 2012]. This is given by equation 2.13.

$$t_n = \pi \left( \frac{k_n}{m_{eff}} - \left( \frac{\gamma}{m_{eff}} \right)^2 \right)^{-\frac{1}{2}} \approx 0.007s \quad (2.13)$$

In equation 2.13  $\gamma$  is the damping coefficient (here,  $\gamma = 0$ ), and  $m_{eff}$  is the effective mass, defined as  $m_{eff} = \frac{m_1 m_2}{m_1 + m_2}$ , where  $m_1$  and  $m_2$  are the masses of the two colliding particles. The effective mass is used to translate binary particle contacts into a center of mass reference frame. Using our lower bound on stiffness (100,000  $N/m$ ) gives us an upper bound on time,  $t_n = 0.007$  seconds. Thus, these interactions are well separated, and we can expect the significant dynamics introduced through deformation to play out while the crests (moves from point B to C in Figure 2). In the next section, we detail the methodology behind this modeling approach, laying the foundation for understanding how these microscale deformation dynamics influence the overall behavior of the granular system.

### 2.3.2 Modeling Deformable Particle Collisions

The equation for springs in series (equation 2.1) can be seen as a simple constitutive equation for a system of springs which succinctly captures an emergent effect. This effect is implicitly captured if each spring is analyzed individually. Similarly, the bulk effect for particles compressing in series can be captured by resolving particle-particle interactions in DEM's such as LAMMPS and MFIX. However, what is missing is a description of the emergent effect, a granular equivalent to equation 2.1; such a description is necessary to complete equation 2.10 and close our analysis of the modified quartier model. This is necessary to be able to say anything prescriptive about phenomenon which emerge from these compressions in series, however, rather than using existing DEM approaches to capture this effect, we choose to model deformable collisions directly.

DEMs typically resolve the contact force between two particles by computing their overlap [Garg et al., 2012]. While this method successfully captures the reduction in work transfer as force propagates down a chain of contacts, it neglects

the subtle gaps and translations that occur during real deformation. Since the contact length scale directly governs the reduction in effective stiffness, its evolution over time is crucial. To fully resolve the contact dynamics arising in simple one-dimensional collisions, it is necessary to model the deformation process explicitly.

To illustrate our approach, we begin with a two-particle example in which one deformable particle slides into another initially at rest. Supplemental materials—including simulation videos for both two- and three-particle systems—provide valuable intuition for the outcomes discussed below, highlighting how contact structures become increasingly complex as additional particles join the compression chain.

Deformation is linked to inter-particle forces through the linear dash-pot model (equation 2.14), where  $k$  represents the material property of stiffness and  $\Delta x$  represents the lengthscale of deformation, which is assumed to be axisymmetric.

$$F = k \cdot \Delta x \rightarrow a = \frac{k}{m} \Delta x \quad (2.14)$$

If  $r_0$  represents the relaxed radius of our particle, then the distance between the two deformed particles can be represented as the sum of the deformed radii. By Newton's third law and equation 2.14 these deformations must be equivalent.

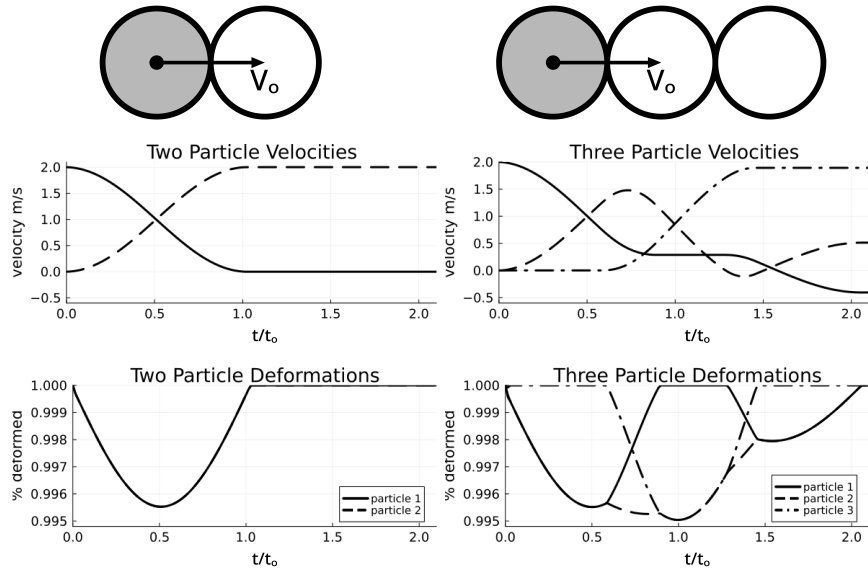
$$x_2 - x_1 = \underbrace{\frac{(2r_0 - \Delta x_1)}{2}}_{\text{particle 1}} + \underbrace{\frac{(2r_0 - \Delta x_2)}{2}}_{\text{particle 2}} = 2r_0 - \Delta x$$

This allows us to write the equations of motion as a function of particle position:

$$a_1 = \frac{k}{m} (2r_0 - (x_2 - x_1)) \quad a_2 = -\frac{k}{m} (2r_0 - (x_2 - x_1)) \quad (2.15)$$

Even for this two-particle system, no analytical solution exists; however, the clear link between acceleration and particle separation makes a numerical solution via the classic 4th-order Runge-Kutta (RK4) method straightforward. The results for this model run for particles of rigidity  $\kappa = 100,000$  (our median rigidity) are summarized in the left hand side of Figure 6, where we recover the expected total exchange of momentum (mobile grain ends with zero velocity).


With two particles, conservation of momentum and energy allow us to solve for the expected deformation and velocities, and we can calculate an error of our numerical approximations which is shown in Figure 7. The error appears to significantly decrease with particle stiffness, becoming less than half-a-percent by



*Figure 6.* Velocity and deformation outputs from the deformable collision model. The left and right side depict two and three particle results respectively.

$k = 10,000N/m$ , which is an order of magnitude smaller than the lowest stiffnesses in our numerical experiments.

Extending the system to three particles, an active particle now collides with a chain of two stationary particles. Here, momentum and energy conservation no longer yield unique solutions, and the deformation becomes asymmetric (it remains axisymmetric but the deformation of two particles in contact no longer has to be equivalent). As illustrated in Figure 6(right), prior to the collision with the third particle, the first two particles exhibit equal deformations. Once the three-particle chain forms, however, the second particle undergoes additional deformation due to its contact with the third, making the second particle the most deformed particle in the chain. To resolve this, the constraint of equal overall deformation, which is used to set up the necessary differential equations for the two particle case (eq 2.15), is replaced by one of equivalent partial deformations. Specifically, the deformation imparted to particle two by particle three must equal the deformation of particle three itself. Consequently, modeling deformable collisions in an  $n$ -particle system becomes an exercise in recovering unique boundary deformations from the particles' center-of-



*Figure 7.* The percent error in deformations for a two particle system as a function of stiffness ( $k$ ). Percent error is calculated between the maximum deformation predicted from conservation equations and the maximum deformation observed in our model. The minimum stiffness used in our simulations is 100,000 N/m.

mass positions. A full description of this procedure is covered in appendix A.1, but we highlight the key points below.

Given the center of masses for an  $n$  particle system, we can guarantee that the two closest centers must be in contact. This provides sufficient information to identify which particles form a compression chain. Once we know which particles participate in a compression chain, the constraint of equal partial compressions can be used to recover the boundaries. The equation for recovering the boundaries of  $n$  particles in a compression chain is given by equation 2.16 and equation 2.17.

$$b_0 = \frac{1}{n} \left( (-1)^n r_0 + \sum_{i=1}^n ((-1)^{i+1} (2(n-i) + 1) x_i) \right) \quad (2.16)$$

$$b_n = x_n + (x_n - b_{n-1}) \quad (2.17)$$

Equation 2.16 finds the leftmost boundary ( $b_0$ ) from the center of masses ( $x_i$ 's) of a  $n$  particle chain, and then equation 2.17 leverages the assumption of axisymmetric deformation to create a recurrence relation for boundaries of particle's in a compression chain (all boundaries can be found from one, but finding the first one is not trivial). Once these boundaries are known, so too are the deformations, and we can implement our RK4 for a system of  $n$  particles.

**In summary, the particle accelerations, positions, and physical constraints encode sufficient information to uniquely reconstruct particle contacts and deformations. Information about particle contacts and**

## deformations can, in turn, be used to create a generalized version of equation 2.15 for an $n$ particle system

Once the deformation and contacts of each particle are recoverable from particle center of mass, an analogous Rk4 model can be used to simulate systems of arbitrary size ( we simulate up to 20 particles). The right hand side of Figure 6 illustrates how quickly the dynamics change with the addition of just one particle. Two characteristic changes in system behavior include (1) an incomplete transfer of momentum, shown in the negative final velocity of the first particle, and (2) an increase in total deformation, and thus energy stored, for the three particle system. From the active particle’s point of view both of these changes are analogous to colliding with one particle of reduced stiffness, manifesting the exact weakening effect we were expecting. Furthermore, while the change in total deformation is small, the stiffness dictates this is a significant difference in energy storage.

Finally, two modifications are incorporated into our deformable particle model before it is used to calculate impulses and ultimately  $\alpha(n)$ : (1) the  $n^{\text{th}}$  particle is fixed in space, representing the substrate’s bottom, and (2) an optional gravitational field in the direction of motion is added. This extended model is capable of capturing both horizontal, frictionless collisions (analogous to an air hockey table) and vertical collisions (such as a grain dropping onto a tower of stationary particles). With this model now fully developed, we next examine its implications for the contact force in the bed-modified Quartier model.

### 2.4 Results

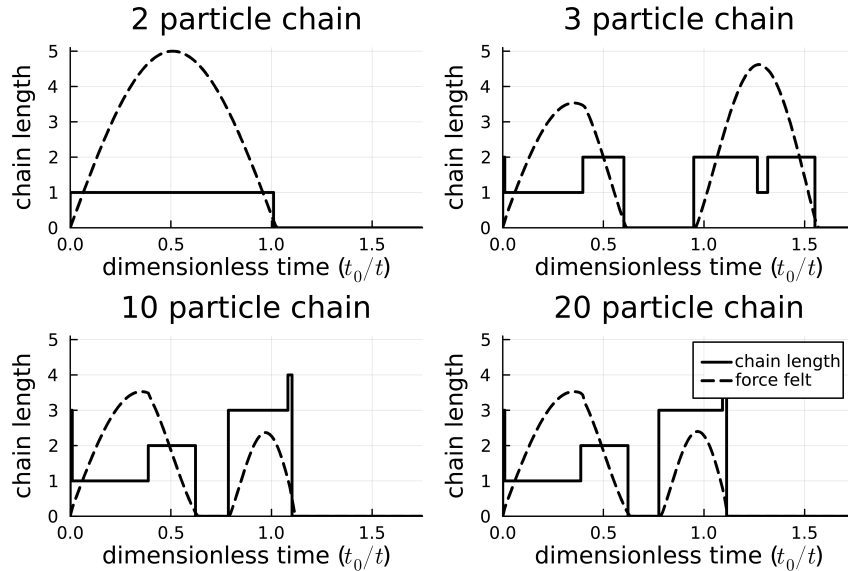
Our results section begins by extracting the impulses experienced by a mobile grain in a one-dimensional deformable collision model and culminates with scaling relationships for thin layer stability observed in real granular systems. To emphasize the relevance of each step toward our broader goal—describing the emergent phenomenon of thin layer stability—we connect every result to specific aspects of Figure 5) This scaffolding demonstrates how our findings progress from the microscopic particle scale to the mesoscale and ultimately to the system scale:

- **Result 1:** By modeling the dynamics of deformable grains, we can capture a modified version of the effective spring constant for massive separable springs, which is necessary for capturing the effect of compressed particles in series. This result is the modeling of the system shown in offset A of Figure 5.

- **Result 2:** From these models, we calculate the impulses felt by a mobile grain colliding into a chain of  $n$  particles. These modified impulses allow us to capture the dynamics at play as the mobile grain leaves the particle well in the modified Quartier model which allows for deformable particles and variable substrate height – the scale at which this phenomenon plays out is represented by the greyed out section in Figure 5
- **Result 3:** Finally, these impulses can be incorporated into the key equations describing the Quartier Model (equation 2.10) to recover the systems critical angle (angle at which perpetual flows is sustained) dependence on system height. This is represented by Figure 5 as a whole.

Our first set of results examines how the contact structures evolve as a mobile grain collides with a static chain of  $n$  particles—and how these evolving structures affect the forces acting on the mobile grain. To emphasize the link to the Quartier model, the particle originally imbued with a velocity will be referred to as the mobile grain, while the remaining static particles form the compressive chain. In these simulations the direction of motion is aligned with the gravitational field, and the last particle is fixed in space. One can therefore interpret these simulations as dropping a grain onto a tower of  $n$  grains, perfectly aligned along center of mass. We ran simulations for  $n$  ranging from one to nineteen (i.e. two to twenty particles total), and highlight results for the two, three, ten, and twenty particle cases in Figure 8. The two and three particle scenarios emphasize the immediate and significant dynamic changes, and the ten and twenty particle scenarios demonstrate the eventual convergence of behavior as the chain length increases.

These plots trace the absolute value of a normalized force felt by the mobile grain and the length of its resulting compression chain through dimensionless time. The force is normalized to the maximum force felt across all 19 simulations, and then scaled up by the contact structure number to plot nicely on the structure plot. This preserves the value of relative comparison between plots (if the peak of one simulation looks twice as large as another simulations peak - it is). Additionally the accelerations due to gravity are subtracted out, to isolate the the forces caused by particle contacts. Time is normalized by the contact duration of the two-particle simulation ( $t_0$ ), ensuring comparability between simulations with different stiffness values. Figure 9 shows that this non-dimensionalization nicely collapses data produced from simulations run

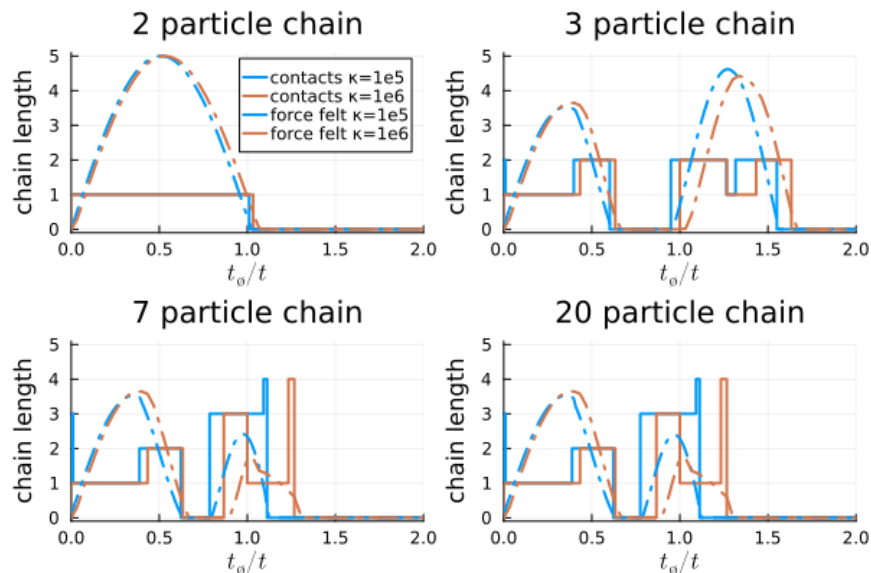


*Figure 8.* The dashed lines trace out the magnitude of forces felt by a mobile grain, under the influence of gravity, as it collides into a vertical chain of particles of various length. Forces are normalized such that proportional comparisons between various chain lengths is accurate. The solid line indicates the number of additional particles currently deforming in series with the mobile grain (chain length). Time is non-dimensionalised such that the two particle interaction is one unit of dimensionless time.

with stiffnesses varying by an order of magnitude. We walk through the three particle simulation of Figure 8 as an example for interpreting these plots.

At the initial timestep, the mobile grain contacts the second particle, which in turn is in contact with the third particle. Initially, only the mobile grain is moving. As time progresses, the first and second particles deform, temporarily breaking contact with the third. These two particles continue to deform and move until a non-dimensional time of  $\sim 0.5$ , at which point the deformed first two particles re-establish contact with the third particle—indicated by the solid line stepping up from a chain length of one to two. Subsequently, as the mobile grain rebounds and the second and third particles continue to deform, there is a transient loss of contact (evident from the drop in both the dashed force curve and the contact line). At one unit of dimensionless time, the decompressing particles re-initiate contact with the mobile grain, with the force signal rising once more and the contact chain length returning to two. These interactions persist until about 1.5 units of dimensionless time, at which point the mobile grain is ejected upward away from the chain

Essential in these dynamics are the alignment between characteristic changes in these force curves with change in contact structure, supporting our premise that the effective force exhibits a series-like dependence on the number of particles in the chain eq 2.1. In our subsequent analysis, these modified impulses will be used to compute the overall impulse for systems of massive, separable springs arranged in series, as anticipated in our introduction.



*Figure 9.* An illustration of how our non-dimensionalization collapses results from simulations run with rigidities separated by an order of magnitude.

Our second result concerns the impulse extracted from our simulations. Once again, these results are expressed in a dimensionless quantity of normalized impulse  $\alpha(n)$  (eq 2.10), which is expressed as a function of chainlength  $n$  to highlight its dependence on the number of particles in the compressive chain. Similar to our dimensionless time, this dimensionless impulse is normalized to the expected impulse in the two particle scenarios, the null hypothesis where relevant interactions occur solely between direct neighbors. To demonstrate the generality of these results, we calculated  $\alpha(n)$  for six sets of simulations with rigidity spanning three orders of magnitude ( $10^4$  to  $10^6$ ) and with two different initial velocities for the mobile grain (2 and 4 m/s), (Figure 10).

A key challenge in extracting the impulse from these simulations is determining the appropriate timescale for integration. (For a comprehensive discussion, see




Figure 10. Blue, green, and red curves show impulse calculations for simulations run at rigidities of  $10^4$ ,  $10^5$ , and  $10^6$  respectively. Lighter and darker shades of each color represent initial velocities of 2 m/s and 4 m/s respectively.

Appendix A.2.) In brief, the impulse is computed over a timescale defined by the maximum compression of the particle chain.

As anticipated, the impulses decrease asymptotically with chain length. This asymptotic behavior arises because small system sizes constrain the formation of compressive chains. Beyond approximately ten particle contacts, the chain length plateaus, and further increases in system size have little effect on the impulse. The x-axis is labeled as non-dimensional height, to emphasize the link between the impulse calculated for an  $n$ -chain one dimensional system, and the capacity for a modified Quartier model of height  $n$  to resist the driving force of gravity. In other words, we are seeing the asymptotic weakening of a system’s geometric friction with increased system size.

Our third result provides the critical link between the microscale dynamics of particle deformation and the emergent macroscale phenomenon of thin layer stability. Using our formal analysis of the modified Quartier model, we compute the system’s critical angle—defined in Equation 2.6—as a function of the modified shock term (Equation 2.10). Remarkably, this analysis recovers the experimentally observed scaling behavior of thin layer stability, as illustrated in Figure 11.

**These findings indicate that the physical mechanism of compression chains is both necessary and sufficient to explain the emergence of thin layer stability in frictionless, monodisperse, inclined granular flows.**





Figure 11. Circles are our analytical results for critical angle as a function of nondimensionalized height (eq 2.6). Triangles and dashed line are Perrin’s experimental results and line of best fit

Furthermore, the fundamental nature of the granular system analyzed here suggests that compression chains play a pivotal role in setting mesoscale dynamics in any granular system where the timescale for compression in series is shorter than that for particle rearrangement. Thin layer stability thus serves as a robust, testable framework for understanding how even subtle softness in nominally rigid grains can lead to significant macroscopic effects. Having demonstrated the existence and scaling of this phenomenon, we now consider its broader implications, extending these insights to more complex and less idealized granular systems.

## 2.5 Discussion

Our findings reveal the significant role of softness in granular systems, particularly through the mechanism of compression chains. By allowing for particle deformation, we have demonstrated that the resulting compression in series influences local coefficients of effective friction, directly affecting the stability of granular flows. However, in most granular systems, compression in series will not manifest in such ideal configurations.

Figure 12 illustrates the contact network for a two-dimensional packing of polydisperse disks as shown in Papadopoulos et al. [2018]. As with any complex contact structure, they are no longer clean, straight lines, but these structures still facilitate compression in series. In lieu of performing a full decomposition of contact forces based off of interior angles formed from the contact network, as



*Figure 12.* Visualization of contact network formed from two-dimensional polydisperse disks as shown in Papadopoulos et al. [2018].

would be required by our original approach, we take a simpler first order approach towards describing how softness effects the energy storage capacity of these contact structures: a behavior which provides a natural bridge for the gap between microscopic particle interactions and macroscopic flow behaviors.

Under the linear-dashpot model, a particle with rigidity,  $k$  subject to force  $F$ , will be compressed by some amount  $c_1 = \Delta x$ . Two of these particles compressed in series have an effective rigidity  $k' = k/2$  (per linear-dashpot assumption), and if subject to the same force will experience twice the compression,  $c_2 = 2c_1$ , as shown below:

$$F = k' \cdot c_2 = \frac{k}{2} \cdot 2c_1 = k \cdot c_1$$

This also means the two particle system will store twice the energy, because the potential energy only depends linearly on  $k$  yet quadratically on  $\Delta x$ . This is shown below where the subscript indicates the systems particle number, and each system is subject to the same force. Thus,  $U_2$  represents the potential energy of a two particle system.

$$U_1 = \frac{1}{2} k c_1^2$$

$$U_2 = \frac{1}{2}k'c_2^2 = \frac{1}{2}\left(\frac{1}{2}k\right)(2c_1)^2 = \frac{1}{4} \cdot 4c_1^2 = 2U_1$$

In general we have  $U_n = nU_1$ . Not surprisingly, the more particles which are able to compress in series, the more efficient that contact structure will be at storing potential energies. This is verified in our numerical models and can be seen by comparing the bottom plots in Figure 6. The three particle system achieves a larger total deformation. This increased capacity to store potential energy can also be viewed as an increased effectiveness at damping velocity fluctuations, an effect which grows like the square-root of the compression chain lengthscale. This is seen in the kinetic energy which can be accommodated by a spring system of length  $n$  (equation 2.18)

$$KE = \frac{1}{2}mv^2 = U_n = nU_1 = n\left(\frac{1}{2}k \cdot c_n\right) \rightarrow v \propto \sqrt{n} \quad (2.18)$$

Revisiting Figure 12, the variations in contact geometry and structure can now be linked to a rich field of information describing local spatial variations in ability to damp velocity fluctuations. Put another way, compression through these contact structures introduce a non-local lengthscale (length of compression) which is determining the diffusion of granular temperature (a description of particle velocity fluctuations). Within the context of our original thin layer stability argument, the lengthscales in the contact network are controlled by the height of the fixed bed. As the lengthscale of the fixed bed increases so too does damping efficiency due to deformability (equation 2.18), which leads to a strong reduction in the contact structures ability to impose geometric friction. As a quick check we can see how this general energy approach could similarly (albeit less explicitly) capture the emergent behavior of thin layer stability. By modifying the velocity loss between energy traps to have the same proportionality as exhibited in equation 2.18, we can produce Figure 13, which recovers thin layer stability up to a lengthscale of approximately ten particles. We replace the previously derived  $\alpha(n)$  scaling—obtained through detailed modeling of deformable particle collisions—with the simpler  $\sqrt{n}$  proportionality revealed by our energy analysis. In effect, the original Quartier model can be modified by incorporating a  $\sqrt{n}$  dependence to reproduce the scaling shown in Figure 13. While the asymptotic behavior is not recovered, our results do not show any particle contacts over ten-particles forming, even for the twenty particle chain. Furthermore this ten-particle lengthscale shows up frequently in the literature as a way to delimit the

mesoscale [Zhang et al., 2014]. In this manner, the asymptotic behavior would be set by the fabrics inability to form many strong contacts longer than 10 particles.

These results underscore the necessity of incorporating deformability into models of granular rheology. By doing so, we can begin to move beyond phenomenological descriptions and towards a more precise understanding of how microscale behavior constructs mesoscale and non-local effects.

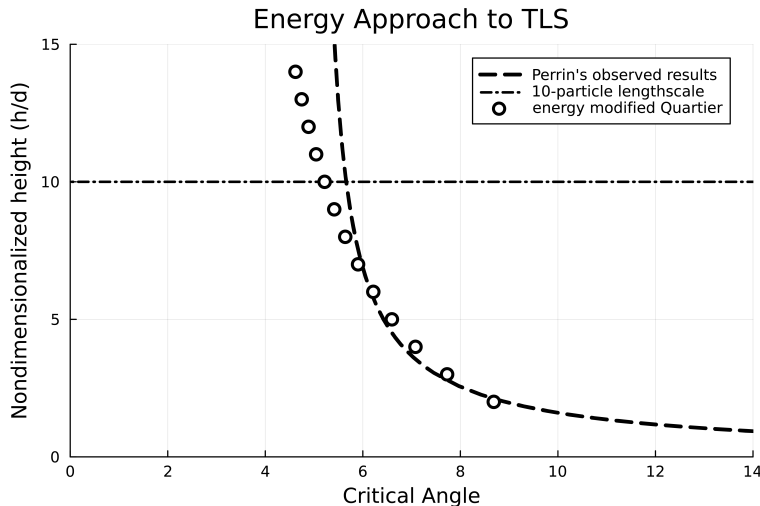


Figure 13. Scaling of critical angle produced by modifying the effective shock (eq2.10) according to the results of our energy budget analysis. The horizontal dashed line represents the 10-particle lengthscale cut off (a common threshold for mesoscale).

## 2.6 Conclusion

By removing the constraint of infinite rigidity, we create space for a mechanism to manifest which uses the compounding effect of compression through chains of particles to modify the effective rigidity in the neighborhoods of its terminal points. While forces ranging an order of magnitude produce deformational differences less than 0.000001% in particles with rigidity of  $10^6$ , allowing for deformation and the formation of compression chains introduces a dynamic which can nearly halve the impulse generated from contact forces (Figure 10). These contact forces, as emphasized in the Quartier model, can play an important role in determining the geometric friction of the system. Accounting for this phenomenon in the Quartier model allows for the construction of a prescriptive description of the critical stresses in inclined, planar, and frictionless granular flow which recovers the hereto unexplained scale-dependent and non-local dependence on system height (thin layer stability). More generally, the number of simultaneously compressed particles introduces an important

new lengthscale (attached to an understood kinematic effect), which links a mesoscale structure to local rheologic ability to dampen velocity fluctuations. This elucidates the natural role of deformability as a bridge between microscopic particle interactions and macroscopic flow behaviors. The propensity for a system to form compression chains and their respective efficacy will be set by properties of the contact network, opening up many fascinating paths of inquiry into the geometry of structures such as Figure 12 . Ultimately, we wish to expose the richness of new interactions which can only arise from relaxing the assumptions of infinite rigidity and to motivate the necessity of grappling with deformation directly if we want to move beyond phenomenological descriptions for mesoscale and non-local effects.

## CHAPTER III

### COMPLETENESS OF THE VOLCANIC RECORD

#### 3.1 Intro

Volcanoes are among the most captivating and dangerous features of our planet. As surface expressions of deep-rooted Earth systems, they provide rare insight into Earth’s interior and interactions with the surface. Volcanic vents mark the interface between surface and subsurface, connecting to dynamic magma transport systems that link deep geologic processes to the surface. Unsurprisingly, these vents are also the sites of some of the planet’s most extreme and unpredictable natural hazards [Ewart, 2005]. Tephra, lava flows, and other volcanic deposits provide a uniquely tangible record of interior Earth systems, mapping the history of volcanic systems on the Earth’s surface.

This physical record, imperfect though it may be, forms the foundation of nearly all modern efforts to understand volcanic processes over long timescales. Thanks to decades of careful cataloging and synthesis, most notably through the Smithsonian Institution’s Global Volcanism Program (GVP), we now possess comprehensive eruption chronologies for hundreds of volcanoes around the world. And yet, even this rich compilation captures only a fraction of what has actually occurred. Estimates suggest that a complete volcanic record should contain  $\sim 40$  million events [Papale et al., 2021]. This incompleteness is a fundamental limitation that hampers our ability to make confident, statistically rigorous statements about volcanic behavior, especially at local and regional scales.

Traditionally, the limitations of sparse and incomplete volcanic records have been addressed by aggregating data spatially, and analyzing the global eruption record as a single, composite system [Rougier et al., 2018, Papale, 2018]. While this is a natural and often useful first step, it risks obscuring local heterogeneity and complicates comparison across different volcanic contexts. Efforts to address volcanism at regional or individual-volcano scales [Coles and Sparks, 2006, Wang et al., 2019] frequently run into a more fundamental obstacle: many of the best-documented volcanoes do not conform to idealized statistical models such as the Poisson process [Bebbington and Marzocchi, 2011, Guttorp and Thompson, 1991, Kiyosugi et al., 2015], which is often the least demanding and most commonly assumed null hypothesis [Cox and Isham, 1980]. In many cases, the available data are sufficient to reject Poisson behavior, but insufficient to move beyond it. This work addresses that gap by opening a new

inferential space through temporal aggregation. By varying the timescale over which events are grouped, we introduce a diagnostic parameter, analogous to the bias-variance tradeoff in statistical learning, that allows us to balance descriptive richness with statistical reliability. Crucially, this parameter reveals the system’s natural information resolution and identifies when a record contains too little data to support any statistical conclusion. The result is a framework that respects the limitations of individual records while extracting as much inferential power as they can support.

Though developed in the context of volcanology, our approach is deliberately general. It is designed for any process where data are incomplete, uncertain, or irregular, where statistical naivety is not an oversight but a reality. By systematically escalating the complexity of the questions we ask, and allowing the data to dictate the scale at which those questions are answerable, we present a pathway from sparse observations to strong inference. Furthermore, progression along this pathway takes place on a scale of information quality, which provides a dimension in which eruptive properties of tremendously different scales can be compared, from geysers to large igneous provinces.

The remainder of the paper is structured as follows. We begin with a brief overview of prior work on the volcanic record and its interpretation. We then outline our methodological framework, highlighting the ways it diverges from traditional approaches. We present a set of preliminary results across a range of volcanic systems, and conclude with a discussion of implications, limitations, and future directions.

### **3.2 Background**

The development of statistical approaches to volcanic systems closely parallels the historical progression of seismology. Early seismological efforts focused on probabilistically describing earthquake occurrences, eventually evolving into magnitude-frequency relationships such as the Gutenberg-Richter law, and then into practical challenges like record completeness [Wiemer and Wyss, 2000]. Similarly, volcanic systems have been studied through a comparable framework: temporal process characterization [Coles and Sparks, 2006, Guttorp and Thompson, 1991], magnitude-frequency modeling [Rougier et al., 2018, Papale et al., 2021], and eventually, efforts to address record completeness and localized behavior [Marzocchi and Zaccarelli, 2006, Wang et al., 2019].

A central challenge, however, arises from the incompleteness of the geologic record, particularly the underrepresentation of low-magnitude or older eruptions. In

seismology, the widespread and frequent nature of earthquakes, along with dense, real-time monitoring networks, has made data incompleteness a largely historical concern. By contrast, volcanic eruptions are rarer and less uniformly observed, making gaps in the record a persistent obstacle for volcanology. Coles and Sparks [2006] addressed this issue by thresholding eruption catalogs by magnitude. This helped mitigate some of the preservation bias affecting smaller eruptions while maintaining Poisson-like temporal properties. Even so, non-poissonian behavior was also observed in several well documented records. Papale et al. [2021] showed the global record behaves poissonian and extended the work to a global scale using the LaMEVE database, introducing a statistically rigorous magnitude-frequency framework. However, LaMEVE’s focus on large-magnitude events and Papale’s reliance on global spatial aggregation obscure the variability essential for understanding regional behavior. Another improvement made by this paper was the inclusion of uncertainty quantifications, something noticeably absent from most of the literature until this point.

Other efforts, most notably the Bayesian framework of Wang et al. [2019], have advanced uncertainty quantification by producing full posterior predictive distributions of global eruption recurrence. While this method complements Papale’s global framework and re-affirms Poissonian behavior for global volcanism, it falls short in providing predictive descriptions for regional volcanism. This is primarily because many regional eruption records remain too incomplete to provide strong likelihood-based constraints, often forcing the posterior to rely heavily on prior assumptions. Furthermore, even reasonably complete records frequently exhibit clear departures from Poisson-process timing, making it difficult to universally justify Poisson-based priors. Unlike seismology, where assumptions about underlying behavior can help identify completeness, in volcanology, a lack of clarity regarding baseline behavior makes this difficult. Consequently, numerous individual-volcano datasets struggle to yield reliable inference under traditional frameworks, impeding initial temporal characterization and subsequent magnitude-frequency and forecasting analyses.

One critical advantage of the present work is its focus on revisiting the temporal characterisation using data from the Smithsonian’s Global Volcanism Program (GVP). An excellent resource, recently summarized by Andrews et al. [2025]. The GVP’s Volcanoes of the World (VOTW) database is the most comprehensive resource available for volcanic eruption records, covering all known Holocene and many

Pleistocene volcanoes. Unlike LaMEVE, the VOTW database does not threshold by eruption magnitude, instead providing an extensible, general-purpose catalog. It includes eruptive histories, tectonic settings, narrative activity summaries, and links to physical rock samples in the Smithsonian’s National Rock and Ore Collection. Since 1968, the GVP has also maintained weekly and monthly bulletins summarizing volcanic activity in near real-time. Despite its strengths, incompleteness remains a concern in VOTW: the apparent increase in eruption rates over the past 200 years is largely an artifact of observation bias, and as Andrews et al. [2025] notes, up to 90% of eruptions over the past two millennia may be missing. Nevertheless, the GVP dataset reveals robust statistical signals, including bimodal eruption durations and multi-modal recurrence intervals, which suggest that physically meaningful structure is present despite noise and gaps.

These advances, both in cataloging and in statistical modeling, highlight real progress, but also underscore the limitations that remain. Chief among these is a mismatch between model complexity and data quality: many eruption records are sufficient to falsify simple models like the Poisson process, but too sparse to support more sophisticated alternatives. The central premise of this chapter is to work within these limitations rather than around them. By adopting a temporal aggregation strategy that lets the data dictate the appropriate analytical resolution, we sidestep the need for unrealistic assumptions and instead recover each record’s informational capacity. This allows us to escalate the complexity of our questions only when the data can support it, yielding a framework that is both respectful of data limitations and capable of producing strong, comparative inference.

### **3.3 Methods**

This chapter introduces a statistical framework for extracting meaningful inference from sparse and uncertain eruption records. The approach centers on temporal aggregation, which transforms a record of exact eruption times into a binary sequence indicating whether an event occurred within each time bin. This transformation introduces a critical degree of freedom: the timescale of discretization, which governs the trade-off between bias and variance. Our method leverages this timescale to identify the portions of the record that support stable inference and define the appropriate resolution at which to analyze each segment. Two key tools underpin the process: the Freedman-Diaconis (FD) bin width [Freedman and Diaconis, 1981], used to determine optimal discretization and a derived equivalent Poisson process,

which enables standardized comparison of eruption intensity inferential power across binning scales.

This section describes how we:

1. select stable portions of each record,
2. quantify information limits using discretization and Poisson intensity,
3. asses structure and distinguishability, and
4. enable rigorous comparison across volcanoes, even when records vary in duration and quality.

### 3.3.1 Selecting Informative Sections of the Record

Before analyzing temporal structure, we must isolate the portion of each record suitable for statistical inference. This requires that the recording quality be approximately consistent. Uncertainties such as dating errors or preservation bias may be present, but must not vary unpredictably within the segment. Because temporal resolution constrains the granularity of any statistical inference, we begin by examining how recurrence intervals can be meaningfully aggregated using the Freedman–Diaconis rule.

To identify such stable segments, we examine the behavior of inter-arrival times (recurrence intervals) using a Freedman–Diaconis bin width applied to a histogram of these values. The Freedman–Diaconis rule selects a bin width,  $h$ , which is equal to twice the inter-quartile range,  $IQR$ , scaled by the cube root of the number of observations,  $n$  (Equation 3.1). This choice minimizes the mean integrated squared error between the histogram and the underlying probability density, providing a principled balance between resolution and reliability. Crucially, the Freedman–Diaconis bin width adapts to the data without assuming a specific distribution form, making it well-suited to uncertain natural records Freedman and Diaconis [1981].

$$h = 2 \frac{IQR(x)}{\sqrt[3]{n}} \tag{3.1}$$

While the specific shape of the underlying distribution need not be known, the Freedman–Diaconis estimator does rely on several general mathematical assumptions. If  $f$  is the probability distribution defined over some interval  $I$  (which may be finite or infinite), then it must satisfy the following:

1.  $f \in L_2$ , and  $f$  is absolutely continuous on  $I$  with an almost everywhere derivative  $f'$
2.  $f' \in L_2$ , and  $f'$  is absolutely continuous on  $I$  with an almost everywhere derivative  $f''$
3. The data are drawn from independent realizations of the same distribution  $f$

Assumption 3 warrants particular scrutiny in our context. It implies both independence and stationarity, two conditions that are unlikely to hold for volcanic eruption records. Inter-arrival may be temporally correlated, and the underlying distribution governing those times may change over the span of the record. Consequently, volcanic eruptions cannot be assumed to be independent and identically distributed (i.i.d.).

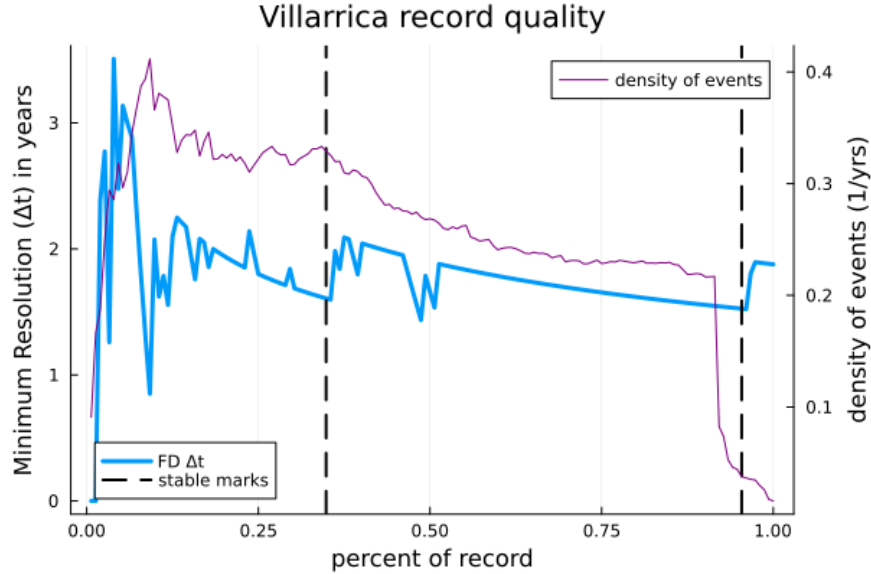
To address this, we apply the Freedman–Diaconis rule not just to the full record, but to progressively expanding, contiguous subsets of the data. If the inter-arrival distribution is highly variable or structurally changing, the resulting bin width  $h$  may behave erratically as more data are added—becoming spiky, unstable, or decreasing too quickly. In contrast, if  $h$  plateaus or decreases smoothly over a stretch of increasing  $n$ , it suggests that the variability is being averaged out, and that the process may be treated as approximately i.i.d. at that scale. In these regions, an optimal discretization emerges organically from the data and reflects the resolution at which short-term dependence and nonstationary behavior are effectively smoothed over.

To formalize stability, we apply the following criteria:

1. **Rolling Stability Threshold:** Relative change in FD bin width remains below 5% across a moving window.
2. **Minimum Segment Length:** Segment must exceed a minimum duration (e.g., 5 units) to avoid short-term artifacts.
3. **Look-Ahead Rule:** Segment is accepted only if no new stable regions emerge shortly afterward.
4. **Termination Rule:** If the record ends during a stable regime, that segment is retained as valid.

5. **Final Index Rule:** The last valid index marks the boundary of the selected segment.

An example of this stable-segment identification is shown in figure 14, with more examples presented as a result in figure 19.



*Figure 14.* The x-axis shows the percentage of the record included, expressed as a fraction of total events. The left y-axis corresponds to the Freedman–Diaconis bin width ( $h$ ) of inter-arrival times, shown in blue. The right y-axis corresponds to the estimated event density, shown in purple. Vertical black lines indicate the locations where stable regions are detected.

Though defining stability involves some subjectivity, we minimize its effects by applying a uniform set of thresholds and rules to all records. For the primary analyses presented here, we focus on the first stable segment detected in each record. This approach emphasizes portions of the record that are likely to reflect modern recording standards and minimizes the influence of long-term preservation biases. Later in the methodology, we describe how to incorporate additional stable segments further back in time. However, such comparisons require relaxing the helpful (but fragile) assumption that longer observation windows inherently reduce variance.

With stable intervals now identified, we turn next to the question of how to characterize the temporal structure of eruptions within these segments.

### 3.3.2 Analyzing a portion: discretization and limits of information

Within each stable segment, we characterize temporal structure using a discretized representation: a binary time series that indicates whether at least one eruption occurred within each time bin. This transformation, reducing a list of continuous event times to a sequence of 0s and 1s, enables statistical modeling with fewer assumptions, but it comes at a cost. By discarding the precise timing of individual eruptions within a bin, we lose temporal resolution and potentially obscure fine-scale dynamics.

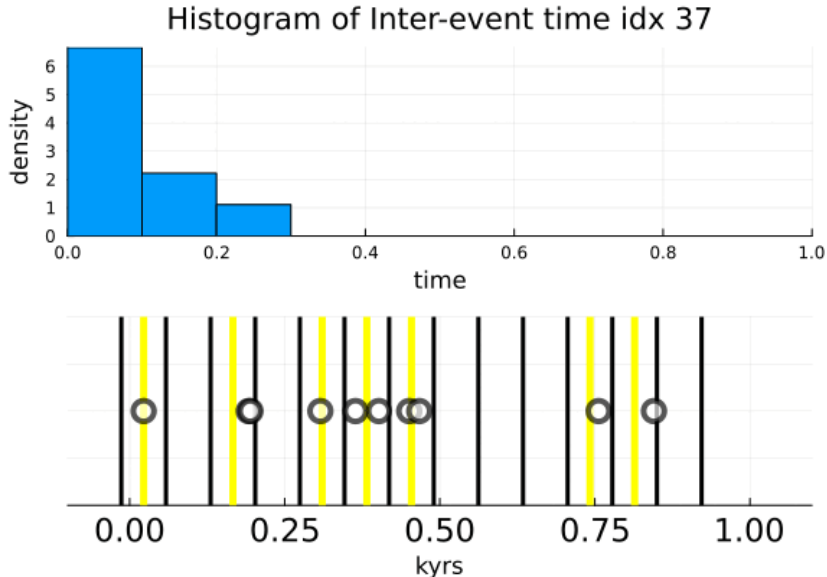
This trade-off reflects the classic statistical tension between bias and variance. At one extreme, analyzing the record with perfect temporal precision minimizes bias but produces high-variance estimates, especially when data are sparse. At the other, very coarse binning reduces variance but inflates bias by collapsing meaningful distinctions between events. The optimal resolution lies between these extremes and depends on the structure and quality of the data.

Once again, the Freedman–Diaconis bin width provides a principled guide. When applied to a histogram of inter-arrival times, it yields a data-driven estimate of the minimum bin width at which the distribution can be reliably resolved. We treat this value as a lower bound on meaningful discretization: any bin size smaller than this threshold cannot distinguish true structure from sampling noise. It sets the finest temporal resolution at which inference is still supported.

To construct the discretized record, we begin by computing a histogram of interarrival times and applying the Freedman–Diaconis rule to determine the bin width  $h$ . This value is then used as the time step  $\Delta t$  for transforming the original point process into a binary time series. An example of this transformation is shown in Figure 15.

At this scale, we reduce variance by aggregating over uncertainties in timing, allowing a wide range of fine-scale processes—including those with subtly different dynamics—to appear statistically similar when viewed at this coarser resolution. The more we aggregate, the more processes become indistinguishable, and thus the less specific information the discretized record contains about the underlying dynamics, but also the less information required to identify structures. What remains is to identify a resolution at which the system is drained of all inferential power

Before we can define such a threshold, however, we must understand how aggregation affects another core property of the system: the intensity of volcanic



*Figure 15.* Top panel: Histogram of interarrival times for events marked in the bottom panel. Bin width ( $\approx 0.1$  time units) is determined using the Freedman–Diaconis rule to discretize the record. Bottom panel: Discretized event timeline; active bins (containing  $\geq 1$  event) are marked with yellow lines, while empty bins remain unshaded.

activity. In the next section, we develop a method for comparing eruption intensity across records that may have been discretized at different temporal resolutions—and in doing so, we uncover a principled way to quantify when a system is drained of all inferential power.

### 3.3.3 Quantifying Eruption Intensity via an Equivalent Poisson Process

Discretized records allow us to measure the proportion of active bins, denoted  $p = \frac{k}{n}$ , where  $k$  is the number of bins with eruptions and  $n$  is the total number of bins set by the length of the record  $T$  and aggregated temporal bin width  $\Delta t = h$ . However, because longer bins are more likely to contain at least one eruption  $p$  is not directly comparable across different bin widths. To standardize comparisons we can transform each  $(p, \Delta t)$  into a continuous time Poisson process with intensity  $\lambda$  using the following equation:

$$\lambda = -\frac{\ln(1-p)}{\Delta t} \quad (3.2)$$

Importantly we do not assume the real process is Poisson. However, when our goal is simply to quantify intensity, rather than to capture temporal structure, then

for any fixed  $\Delta t$  the only pertinent information is the number of active bins. Now the discretized record can be viewed as a sequence of Bernoulli trials, with probability of success  $p = \frac{\text{active bins}}{\text{number of bins}}$ . Under the assumption of a memoryless Poisson process, the probability that a bin is inactive is equal to the probability that no events occur in a window of duration  $\Delta t$ .

This gives the identity:

$$1 - p = \int_{\Delta t}^{\infty} \lambda e^{-\lambda t} dt = e^{-\lambda \Delta t}$$

When solved for  $\lambda$  we get equation 3.2. For a fixed  $\Delta t$  the function  $p \mapsto \lambda(p)$  is strictly increasing for  $p \in (0, 1)$  which guarantees a unique mapping of possible equivalent Poissons to possible intensity values in our discretized record.

The intensity returned from this equivalence not only allows for a nice comparison between records of different discretizations but it can also make explicit the lack of distinguishability between potential generating sources that comes with increased bin size. We do this by perturbing the estimated probability in the discretized record by some fixed margin (say  $\pm 5\%$ ) and tracking how these small differences in  $p$  propagate into differences in the corresponding  $\lambda$ .

A careful reader (and even a not-so careful one) might wonder why we introduce this equivalent Poisson framework at all. After all, eruption density, the number of events divided by total time, already offers a simple, continuous estimate of activity. But this reasoning misses a crucial point: **our confidence is not in the original event times themselves, but in the aggregated structure derived from them**. The binary time series, though coarser, is a more robust and interpretable object precisely because it smooths over uncertainty in timing and completeness. The equivalent Poisson process should therefore not be seen as an estimate of the true, underlying continuous-time dynamics, but rather as a continuous approximation to the *discretized representation* - the version of the record in which we have the greatest epistemic confidence.

To demonstrate the robustness of this intensity estimation framework, we turn to a synthetic experiment. We simulate one-hundred realizations of a simple process: uniformly scattering 10 events over the unit interval. For each realization, we apply the Freedman–Diaconis rule to determine the appropriate bin width, then compute the proportion of active bins  $p$  and the corresponding equivalent Poisson intensity  $\lambda$ .

To test the sensitivity of the method, we select the two realizations that yield the most extreme values of  $p$ —the highest and the lowest. These represent the most

visually distinct outcomes in the simulation. The results are shown in Figure 16. While most realizations look broadly similar, these two serve as endmembers for assessing how much visual variability translates into inferential distinction.

Several key observations emerge:

- The two outcomes produce notably different Freedman–Diaconis bin widths. This reflects the fact that not all realizations are equally informative: the top example exhibits a narrow range of inter-arrival times and is therefore assigned a coarse binning scale, while the middle example shows more varied inter-arrival structure and receives a finer bin width.
- Despite their visual differences, the estimated Poisson intensities derived from each record are statistically indistinguishable. As shown in the right-hand panel, the confidence intervals around  $\lambda$  overlap substantially, correctly capturing the fact that both realizations are drawn from the same underlying process.
- The method adapts naturally to the informativeness of each sample: more varied sequences yield narrower confidence intervals, while less informative ones reflect their uncertainty through wider intervals.

It is worth emphasizing that this is a worst-case comparison: among 100 realizations, we intentionally selected the two most extreme outcomes. In most cases, the visual and statistical differences would be much smaller. Because this is a synthetic system, we know the true underlying intensity ( $\lambda = 10$ ), which might make these extra steps for comparison seem unnecessary. But in real-world data, where the true process is unknown and the exact event times may be uncertain, our aggregated and discretized representation becomes the more reliable foundation for inference. The equivalent Poisson framework respects this structure: it does not exaggerate superficial differences between records, and it accurately communicates the degree of ambiguity introduced by binning and finite sample size.

While this equivalent Poisson method was developed to estimate intensity, the confidence intervals it returns play a deeper role in assessing how much inferential power remains. We demonstrate how this works in the next section where we turn our attention to structure.

### 3.3.4 Comparing Structure and Inferential Power

To enable meaningful comparisons across eruption records of varying quality, we now turn to the explicit decomposition of bias and variance that underlies the

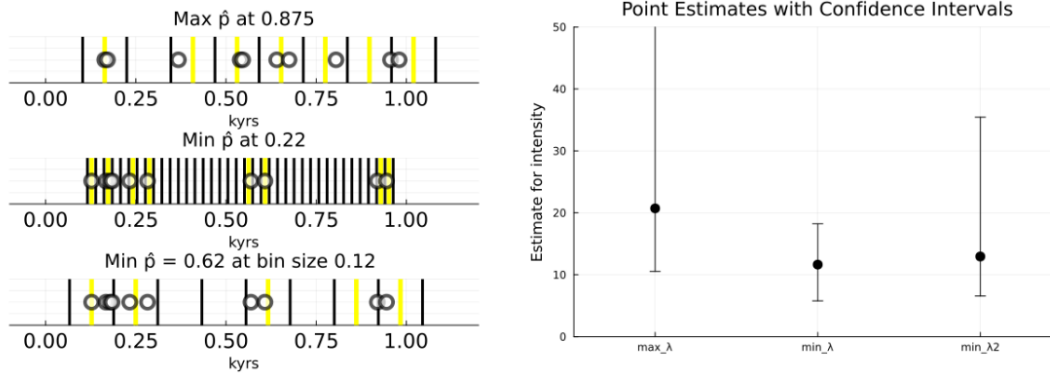


Figure 16. fig:Left-hand-side: The top time series exhibits the highest Freedman Diaconis binwidth out of 100 simulations of  $Uniform(\lambda = 10)$  over the unit interval. Middle represents the smallest binwidth observed from the same simulation set. The bottom record shows the higher resolution record being held to the same standard as the lower. Right-hand-side: From left to right, intensity and confidence interval estimates for top, middle, bottom

Freedman–Diaconis bin width introduced in Equation 3.1. This decomposition, shown in Equation 3.3, originates from Freedman and Diaconis [1981]’s foundational work on histogram-based density estimation:

$$MISE(h) \approx \underbrace{\frac{1}{kh}}_{\text{variance}} + \underbrace{\frac{\gamma}{12}h^2}_{\text{bias}} \quad (3.3)$$

This equation expresses the mean integrated squared error ( $MISE$ ) of a histogram estimator as a function of bin width  $h$ , approximated by the sum of a variance term and a bias term. The variance component,  $\frac{1}{kh}$ , depends inversely on both the number of events  $k$  and the bin width  $h$ , reflecting the familiar increase in variance from undersmoothing. The bias term,  $\frac{\gamma}{12}h^2$ , grows quadratically with bin width and depends on the roughness of the underlying density via the parameter  $\gamma = \int_I f'(x)^2 dx$ . In practice, this dependence on the density’s derivative can be approximated using twice the interquartile range when the true distribution is unknown. As Freedman and Diaconis [1981] note with some humor, this substitution belongs to the realm of mathematical folklore, though more formal numerical justification is available in Scott et al. [1980].

While the previous section demonstrated how to make powerful comparisons of eruption intensity, one might still worry that higher intensities in some records could

arise merely from better resolution (i.e., lower bias) rather than genuine differences in eruptive behavior. To address this, we can perform comparisons at a fixed level of bias across all records which equation 3.3 shows us is achieved by holding bin width  $h$  constant. This ensures that each record is discretized at the coarsest Freedman-Diaconis resolution encountered, standardizing the smoothing effect and preventing overinterpretation of apparent intensity differences that are simply artifacts of finer binning. An added benefit of this approach is that all discretized records will align temporally—i.e., the  $i^{\text{th}}$  bin in each record corresponds to the same time interval—making this strategy ideal for synchronized analyses of eruption rates and for performing correlation analyses between records.

However, if our goal is instead to compare temporal structure, for example, to test for clustering, regularity, or other departures from randomness, then holding the bias fixed is less appropriate. Structural comparisons benefit more from holding the variance term constant across records. As equation 3.3 shows us, we equalize statistical uncertainty by ensuring that  $\frac{1}{kh}$  is the same for all records, where  $k$  is the number of events and  $h$  is the bin width. This allows us to compare patterns on a level playing field: each record is discretized just enough to keep the variance of binned event counts the same, regardless of total length or activity level. A more formal justification for normalizing by variance being the natural choice for comparing structure can be made with the Central Limit Theorem, and is shown in appendix B.

A second essential check concerns inferential power. Even if a record passes the variance- and bias-related thresholds, it may still contain too little information to support meaningful hypothesis testing. We define a minimum threshold for inclusion based on whether the discretized record contains enough statistical power to reject the null hypothesis of a random (i.e., unstructured) process at a significance level of  $\alpha = 0.10$ . For instance, consider the top realization in the synthetic example shown in Figure 16: although it contains eruptions in 9 out of 10 bins, this outcome is indistinguishable from many others that would arise under randomness. In such a case, there is insufficient evidence to conclude anything about structure.

To quantify inferential power, we leverage the confidence intervals derived from the equivalent Poisson formulation introduced earlier. This is best illustrated by example. Suppose we analyze a system with 100 bins, but the estimated probability of a bin being active has a confidence interval of  $\pm 20\%$ . At this resolution the estimated activation probability is effectively constrained to just a few distinguishable values,

sa  $\{0, 0.25, 0.5, 0.75, 1\}$ . While the original outcome space of the gull binary record was  $\Omega = 2^{100}$ , the image of the probability measure for the random variable that  $i^{th}$  bin is active is restricted to one of 5 values  $P\{X_i(\omega : \omega \in \Omega) \mapsto \{0, 0.25, 0.5, 0.75, 1\}$ . Thus the preimage is is much smaller than the original size of  $\Omega$ . This compresses the original inference space: rather than drawing conclusions from the vast space of  $2^{100}$  binary sequences, we are effectively working within a reduced space of just 5 equivalence classes, each defined by a plausible value of pp within the confidence bounds. The system behaves, for inference purposes, as though it were drawn from a set of only a few representative patterns, not the full enumeration of possible binary outcomes.

To asses inferential power in this reduced space we simulate the set of all possible outcomes for  $n$  bernouli trials with probability equal to the estimated  $p$  and we check for pairwise independence at a minimum  $k = 2$  which at least ensures that outcomes like  $1, 1, 1, 1, 1, 0, 0, 0, 0, 0$  or  $1, 0, 1, 0, 1, 0, 1, 0, 1, 0$  are rightfully viewed as unlikely under an assumption of randomness. By enumerating all such outcomes (more details given in appendix B). We find the exact distribution of possible pairwise events, and the CDF gives us the lowest significance we can achieve in testing for random as seen by it's y-intercept. As a demonstration, Figure 17 shows the probability of rejecting randomness across 10 systems, each with  $p = 0.5$  as a function of the effective inference space size  $n$ .

The full methodology for evaluating the inferential power within a discretized record proceeds as follows:

1. Estimate the activation probability  $p$  using the discretized record.
2. Use enumeration (as in Figure ZZZ) to find the minimum inference space size  $nn$  required to reject the null hypothesis of randomness at a desired significance level (e.g.,  $\alpha = 0.05$ ).
3. This threshold defines a minimal confidence interval width of approximately  $\frac{1}{n}$ .
4. Compute the actual confidence interval for pp using the equivalent Poisson formulation.
5. If the width of this interval is smaller than  $\frac{1}{n}$ , the record retains sufficient inferential power to support meaningful hypothesis testing.

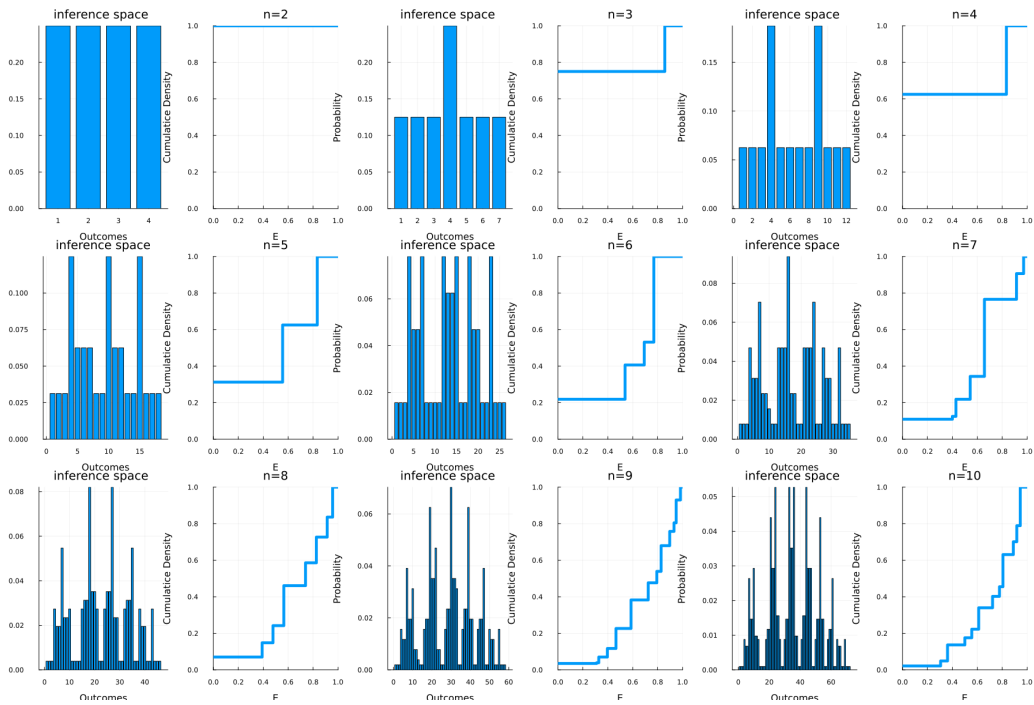


Figure 17. Subplots iterate from reduced inference space of size 2, which results in 4 pairwise outcomes  $\{(0, 0), (0, 1), (1, 0), (1, 1)\}$  to a reduced inference space of 10 which as  $2^{10}$  possible outcomes. The left portion of each subplot shows the full distribution of possible outcomes on each reduced inference space, and the right subplot displays those outcomes distinguishable values as a CDF. The y-intercept on each subplots CDF gives us the lowest possible significance level that can be achieved on the corresponding inference space.

### 3.3.5 Comparing Records of Different Duration

Finally, we address how to compare eruption records with vastly different time spans. In principle, longer records reduce variance in statistical estimates. However, when older eruptions are sparse or poorly constrained, this assumption no longer holds. To expose the role of record duration in our variance term, we decompose it as follows, where  $TT$  represents the total length of the record:

$$\frac{1}{kh} = \frac{\text{number of events}}{\text{number of bins}} \cdot \frac{1}{T}$$

This decomposition shows that variance is inversely proportional to the total observation window  $TT$ , even if event and bin counts are held constant. To avoid spurious comparisons driven by differences in observation length—particularly when longer records are less complete—we remove this  $1/T1/T$  factor and instead hold the bin density constant. Bin density is a dimensionless quantity proportional to variance, defined as the number of events per bin, and it allows comparison without privileging longer records.

Although this adjustment is not tied as directly to the bias–variance decomposition as our earlier procedures, holding bin density constant ensures that the expressive capacity per data point is comparable across records. This final correction allows both recent, densely sampled records and older, sparsely sampled ones to be evaluated on as equal a footing as possible, even when the quality of recording cannot be assumed to be equivalent.

### 3.3.6 Summary of Methods

This methodology outlines a comprehensive framework for extracting reliable inferences from eruption records, even when those records are sparse, incomplete, or unevenly preserved. Beginning with a principled selection of stable record segments using Freedman–Diaconis bin widths, the approach discretizes each sequence into binary time series that retain statistical signal while reducing the demands placed on the data. From these representations, we derive equivalent Poisson processes that quantify eruption intensity in a way that is robust to differing bin sizes and record lengths.

The framework also introduces variance- and bias-normalized comparisons to support structured hypothesis testing. By computing inferential power from a reduced outcome space, we can rigorously assess whether our more reliable (albeit simplified)

representations of uncertain volcanic records contain sufficient information to support meaningful statistical inference. Additionally, adjustments for record duration via bin density ensure that both short, dense records and long, sparse ones can be compared.

While this methodology section lays out the analytical machinery in full, the following results serve to demonstrate its practical implementation. Each step—from discretization to intensity estimation to structure comparison—is illustrated using data from the Smithsonian’s Global Volcanism Program [Global Volcanism Program, 2025], highlighting how the framework executes across diverse data scenarios.

### 3.4 Results

Our results are based on data from the Smithsonian’s Global Volcanism Program (GVP). We include all volcanoes with confirmed Holocene eruptions and more than 50 documented events. For each volcano, eruptions are compiled into a unified dataset with event times rounded to the nearest year. In addition, we incorporate eruption data for Mount Mazama based on Charles Bacon’s geologic reconstruction of its eruptive history. The time-series data from the GVP is shown in figure 18 and amounts to forty-one different volcanos.

The analyses that follow mirror the structure of our methodology and are derived from this compiled dataset. We begin by identifying the most statistically viable segments of each eruption record. We then compare eruption intensities across all volcanoes using the equivalent Poisson formulation. Next, we assess temporal structure, identifying which records retain sufficient inferential power to support meaningful comparisons. Finally, we examine cross-scale comparisons, focusing on three case studies: Villarrica, Sheveluch, and Mount Mazama.

#### 3.4.1 R1: Selecting Statistically Viable Portions of the Record

We apply the Freedman–Diaconis-based record selection method described in Section 3.3.1 to identify the statistically viable portions of each eruption record. Stable regions are defined as those in which the Freedman–Diaconis bin width changes by less than 5% across a rolling window and contain at least 5 events.

Figure 19 summarizes the results for nine volcanoes that span a broad range of eruption behaviors. Each subplot displays the minimum Freedman–Diaconis bin width (blue line, left axis) and the event density (purple line, right axis), both plotted against the percentage of the record’s total length. Dashed black lines indicate the endpoints of selected “stable” segments.

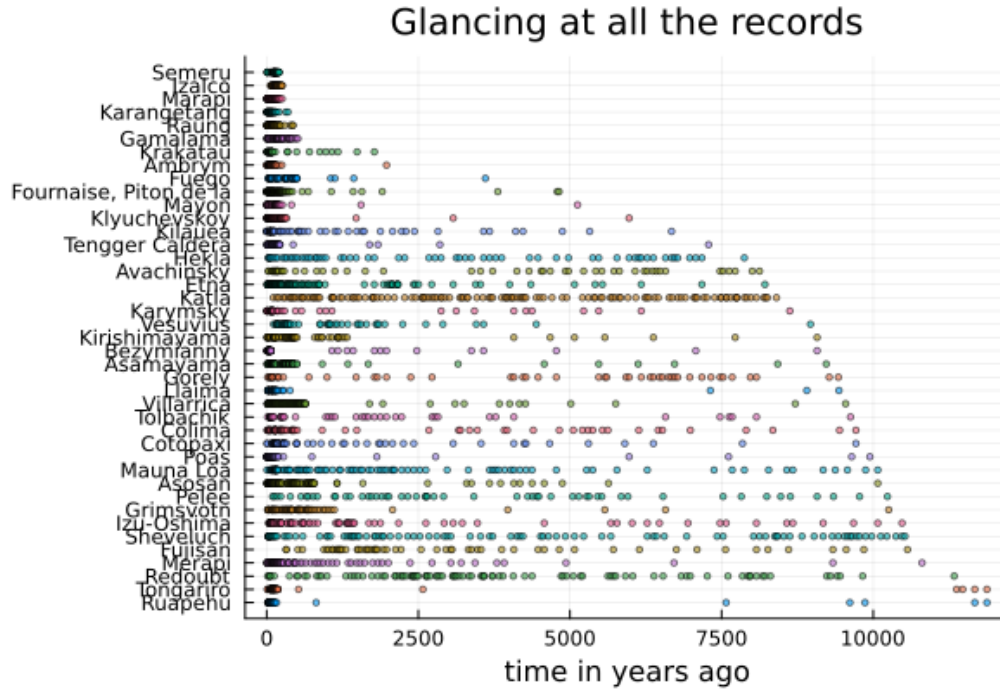


Figure 18. fig: The forty-one Volcanoes in the GVP’s Holocene eruption catalog with more than 50 events cataloged

This figure highlights two key observations. First, the quality and temporal behavior of eruption records vary dramatically across volcanoes. Second, while change points identified through bin width stability often coincide with shifts in event density, this is not always the case—suggesting that density alone is insufficient for identifying viable segments. For example, consider a volcano that produces a dense cluster of eruptions early in its record, followed by a much sparser phase due to an abrupt drop in recording capabilities. Even if events in the second portion are recorded consistently at this new threshold, the average density will continue to decline for some time, masking the regime change. Despite the diversity of records, the stability selector performs well across all cases, generating relatively few false positives. One notable exception is the third stability marker in the subplot for Sheveluch, which is triggered due to that volcano’s Freedman–Diaconis bin width spanning over an order of magnitude. In this case, the 5% stability threshold—calibrated to a global mean—becomes too permissive at the local scale.

We can take the largest stable portion of each record to create a subset of the original data which makes the most use of the data. In this scenario we can’t assume that the recording quality is the same between all the portions, but we can

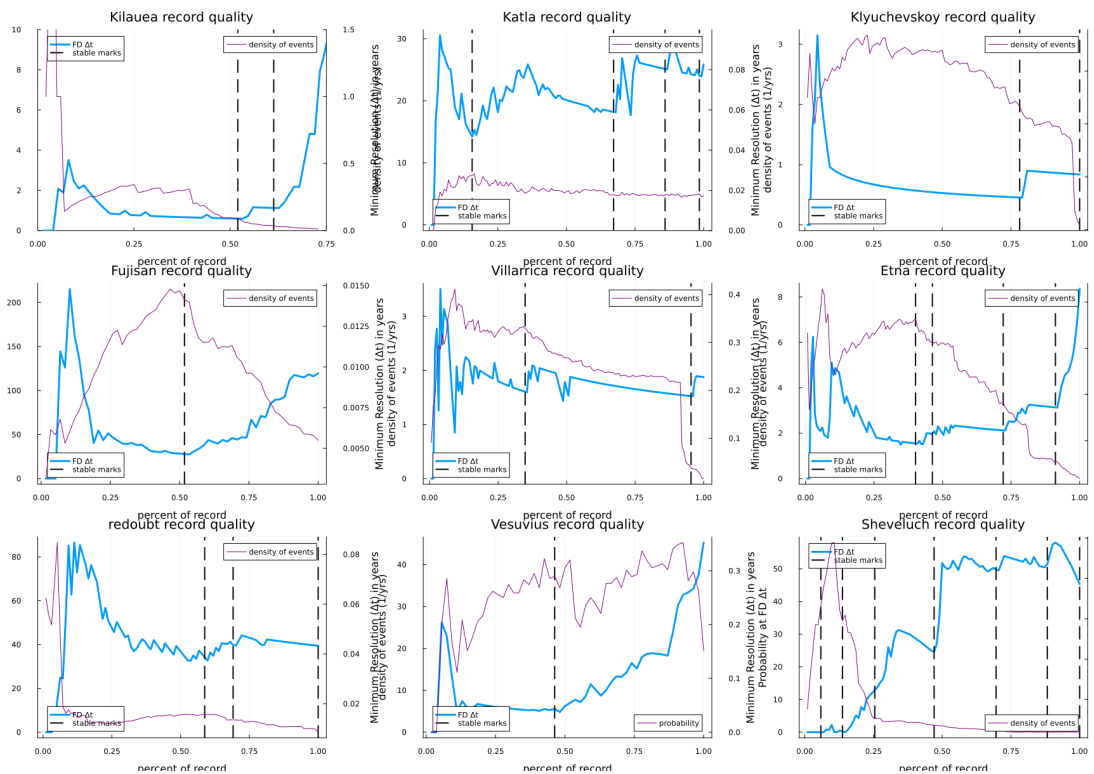


Figure 19. Each subplot contains the minimum FD binwidth plotted in blue and measured on the left axis, the density of events in purple measured on the right axis, all plotted against percent of record on the x axis. The dashed black lines indicate the end of "stable portions" and demark the portions of the volcanic record.

get around that by using the analysis outlined in section 3.3.5. A visualization of this data is plotted in figure 20. Of particular note are the 10 longest lived records which appear to have consistent (not complete) recording going back 10,000 years. It is important to note that many of these volcanoes will have experienced a change of recording quality in their stability analysis, but the record managed to stabilize again , after this change, according to our Freedman-Diaconis measure.

We extract the longest stable portion from each eruption record to generate a reduced dataset that maximizes the use of reliable information. While we cannot assume consistent recording quality across these segments, we account for these differences using the bin-density-based comparison framework outlined in Section 3.3.5 and demonstrated in a later section of the results. A visualization of this subset is shown in Figure 20. Of particular interest are the ten longest-lived records, many of which extend back nearly 10,000 years with reasonably consistent—though not necessarily complete—recording. It is important to note that several of these volcanoes exhibited detectable changes in recording quality during the stability analysis. However, in many cases the record stabilized again after such transitions, as assessed by our Freedman–Diaconis-based selection criteria.

By selecting the first stable portion of each record, we assume these segments most closely reflect modern recording standards—where observational completeness is likely to be highest. Under this assumption, we can conduct our most rigorous statistical analyses, as we treat the recording quality as roughly comparable across records and therefore interpret the length of observation as a meaningful control on variance. These selected portions (shown in Figure 21 form the basis of the following results. It is worth noting that Bezymianny and Krakatau are two volcanoes for which the densest regions of activity do not appear to be properly selected (bottom two plots in figure 21). This shortcoming arises from the Freedman–Diaconis rule’s reliance on the interquartile range (IQR) to approximate the roughness of the underlying density’s derivative. In both cases, several eruptions from much later periods were recorded before their influence was detectable through changes in the IQR. While such pathological cases are expected in a dataset of 41 volcanoes and could be addressed individually, we prioritize maintaining a uniform and transparent methodology in this initial analysis.

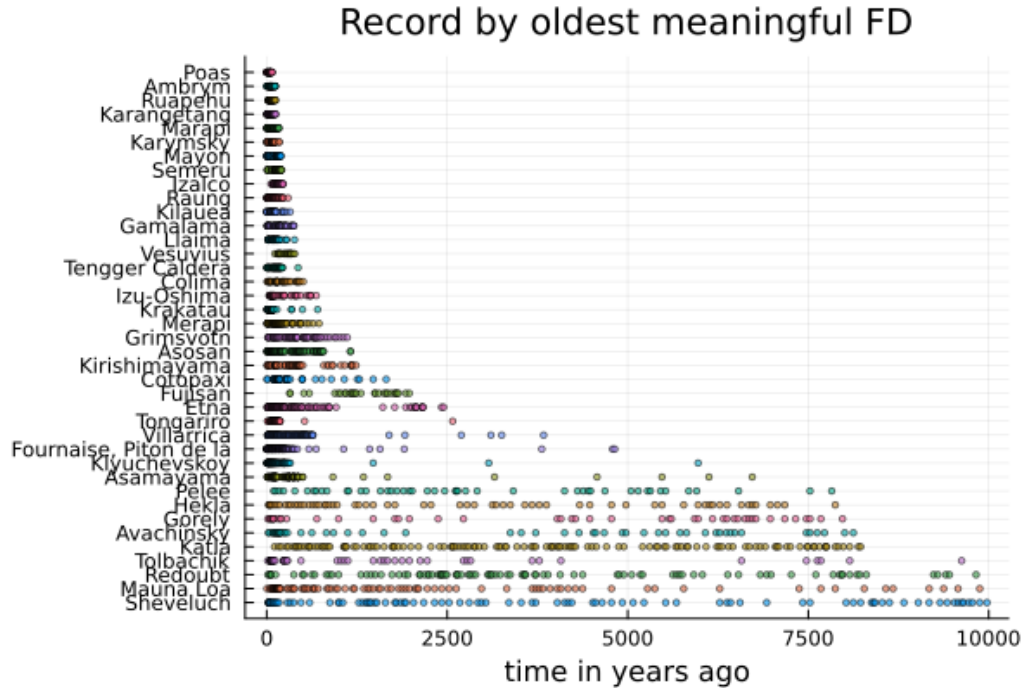


Figure 20. GVP only showing the records included up until the last stable portion ends for each volcano.

### 3.4.2 R2: Comparisons of intensity

We begin by computing the eruption intensity of each volcano individually, using its own optimal Freedman–Diaconis bin width as determined in the previous section. Figure 22(a) displays the density of events (number of eruptions divided by record duration), while Figure 22(b) shows the eruption intensities estimated from the equivalent Poisson process, along with associated uncertainty intervals.

One of the most notable features of Figure 22(b) is the exceptionally high intensity of Piton de la Fournaise. This volcano’s lead likely reflects its extraordinarily detailed observational record, with multiple eruptions per year documented over the past century. However, because our discretization uses a minimum resolution of one year, Fournaise likely fails to reach its true Freedman–Diaconis resolution. To correct this, future versions of the dataset should incorporate sub-annual recurrence times, leveraging the known start and end months of individual eruptions.

Another striking result is the extremely wide confidence interval for Sheveluch, driven by the fact that only six eruptions fall within its selected stable portion. While Sheveluch is a long-lived and well-monitored system, its lack of recent eruptions yields

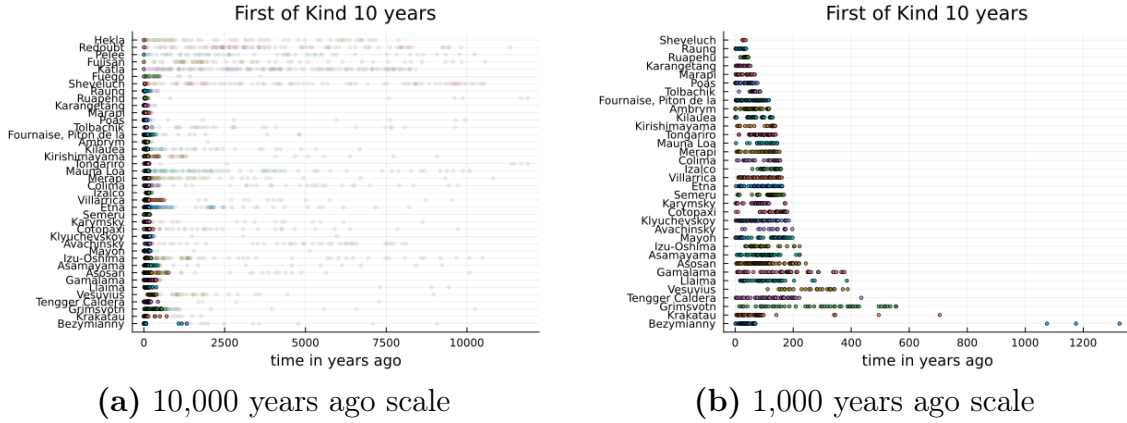


Figure 21. (a) The full record for each volcano is shown at high transparency, while the portion that comprises the first stable portion is shown opaque. (b) The same records but only showing the events included in the first stable portion and zoomed in to show a more representative time scale on the x-axis.

high uncertainty in intensity estimation. So this is the confidence intervals working as expected. Similarly, as noted previously, Bezymianny and Krakatau appear artificially low in intensity due to early outlier events that distort the Freedman–Diaconis binning.

Beyond these edge cases, the consistency across the remaining records is noteworthy. Many volcanoes have similar intensities, despite differences in geography, tectonic setting, and recording history. This re-frames the typical question of “Which volcano is most active?” Instead, a more compelling line of inquiry emerges: might there be a shared physical mechanism setting eruption rates across seemingly disparate systems?

Figure 23(a) shows the intensities measurements when we hold the bias constant to the largest freedman diaconis bin width in the set of modern records in red, contrasted with the original best estimates in black. While the previous intensity estimates are all consistent within their own record this ensures that one record having a higher intensity is not because of some extra structure the lower resolution record misses out on. This also has the added benefit of aligning all the bins temporally as shown in figure (23(b)). This discretized version of the data would be the ideal starting place to begin a correlation analysis between records. However, at this resolution many of the records are mostly active most of the time, illustrating the point that the statistical power of what we can say will be heavily limited by which volcanoes we include. The statistical questions we could ask for this range of volcanoes

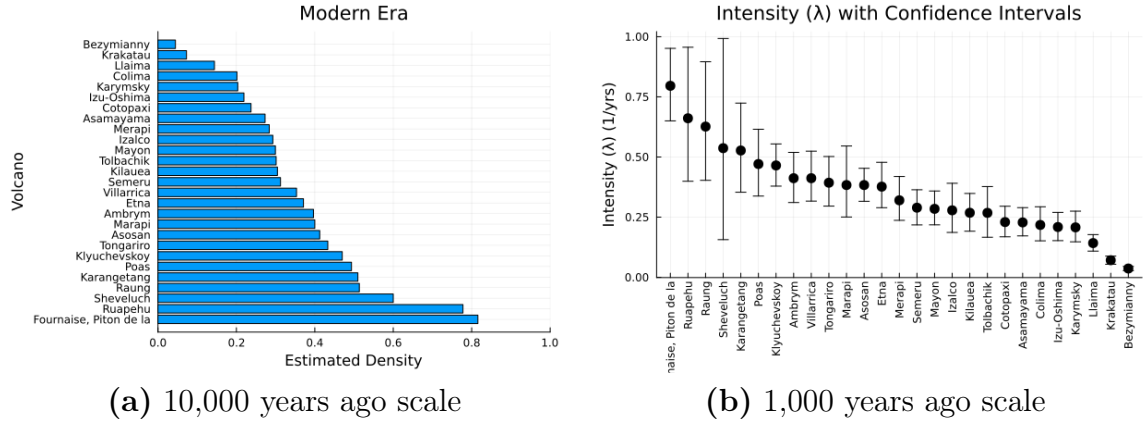


Figure 22. (a) The full record for each volcano is shown at high transparency, while the portion that comprises the first stable portion is shown opaque. (b) The same records but only showing the events included in the first stable portion and zoomed in to show a more representative time scale on the x-axis.

would have to be more focused on the periods of inactivity, or about the intensity in general. However, by narrowing the set of comparable volcanoes to only include those whose Freedman-Diaconis bin widths are smaller than or equal to Kilauea, we greatly increase the information available for comparison. This is illustrated in figure 24 which illustrates that by only include records which have a Freedman Diaconis bin width at least as small as Kilauea our intensity estimates are much closer to the original optimal estimates (figure 23(a) and the discretization picks up on higher resolution periods of activity/empty bins (figure 23(b). This elucidates the importance of matching the type of question ask to the set of volcanoes we hope to compare, but also provides hope in that the 9 volcanoes comparable to Kilauea still offer quite a lot of structure and resolution for analysis.

### 3.4.3 R3: Comparing Intensity Under Matched Bias (Temporal Resolution)

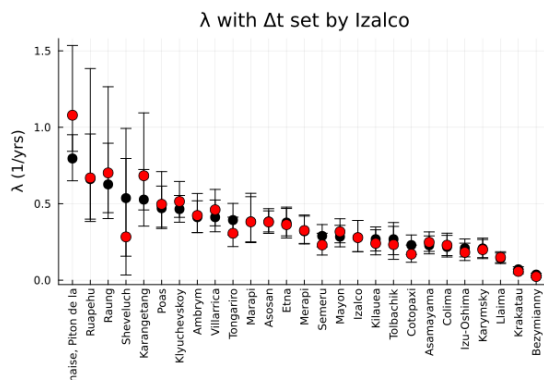
Figure 23(a) compares eruption intensities across volcanoes when the discretization bin width is held constant at the largest Freedman–Diaconis value among the modern records. These uniform-resolution estimates are shown in red and contrasted with each volcano’s individual optimal estimate (in black). While the original estimates are valid within each record’s best-fit resolution, this standardization ensures that observed differences in intensity are not simply artifacts of differing bin sizes. In other words, if one record appears more active under this

coarser resolution, it reflects a genuine difference in event rate rather than a hidden gain in resolution.


An additional benefit of using a shared bin width is that it aligns the temporal bins across all records, as shown in Figure 23(b). This synchronized discretization creates a natural framework for cross-record comparisons, including correlation analyses. However, the tradeoff is evident: at this coarse resolution, many volcanoes appear to be continuously active time for large blocks of time, limiting the discriminatory power of the analysis. Under such conditions, statistical inference would be most informative when focused on rare gaps in activity or broader assessments of overall eruption frequency.

To recover more nuanced temporal structure, we can restrict the comparison to a subset of volcanoes with sufficient data quality to support finer resolution. Figure 24 demonstrates this by narrowing the analysis to only those volcanoes whose Freedman–Diaconis bin widths are less than or equal to that of Kilauea. Within this higher-quality subset, the standardized intensity estimates closely resemble the original optimal values, and the corresponding discretized records capture more detailed patterns of activity, including empty bins and clustered eruptions.

This illustrates an important methodological point: the types of questions we can ask depend critically on the subset of volcanoes included. Coarse-resolution comparisons allow for broader inclusion but limited inference; finer-resolution analyses offer deeper insights, but only for a more selective group. Encouragingly, the nine volcanoes comparable to Kilauea still exhibit rich temporal structure, suggesting that meaningful high-resolution comparisons are both feasible and informative within this group.



(a)  $\lambda$  at Izalco resolution



(b) binwidths aligned in absolute time

Figure 23. (a) Intensity's calculated from the equivalent Poisson when  $\Delta t$  is set to Izalco's resolution for each record (red) compared to the original optimal estimates (black). (b) Black lines represent a bin containing at least one eruption. All bins are aligned in time, such that the  $i^{\text{th}}$  bin in each record corresponds to the exact same interval of time.



### 3.4.4 R4: Comparing Records Under Matched Variance

As outlined in Section 3.3.4, comparing structural features across records is best accomplished by holding the variance term  $\frac{1}{kh}$  constant, where  $k$  is the number of eruptions in the record and  $h$  is the bin width. Since this may require setting  $h$  smaller than a record’s Freedman–Diaconis minimum, we restrict our comparison to records where the adjusted bin width remains at or above the Freedman–Diaconis threshold. We also apply the inferential power check described in Section 3.3.4, ensuring that this re-discretization does not degrade the resolution so severely that the ability to conduct, even simple hypothesis testing is lost.

To mirror the previous analysis, we implement this comparison twice: once using a maximum allowable variance (i.e., the coarsest resolution), and once using a median value to allow more structure to emerge. The maximum variance is initially set by the most extreme records, including Bezymianny and Krakatau, which—despite being included in the modern record subset—have unusually high variance due to how their records interact with the IQR range (as discussed in the first section of the results). These volcanoes likely do not belong in a direct comparison, and their inclusion would render all structure-based inference trivial. To avoid this, we instead set the maximum variance based on the third-highest value in the set, corresponding to Ruang, providing a more balanced benchmark.

The results are shown in Figure 25, which displays histograms of inter-arrival times for the 24 volcanoes whose adjusted bin width meets or exceeds their Freedman–Diaconis minimum. Each subplot’s displays a histogram of inter-arrival times, with the number of bins reflects the resolution—or complexity—of structure we can resolve. Subplots colored orange indicate records that, under this resolution, do not retain enough inferential power to reject a null hypothesis of randomness at a significance level of  $\alpha = 0.1$  (as determined by the methodology in Section 3.3.4).

Following the same logic, Figure 26 displays the results when variance is set to a more selective threshold, specifically, the median variance among all qualifying records, which corresponds to Asosan. Although fewer records meet the Freedman–Diaconis constraint at this higher resolution, those that remain are resolved at greater complexity (i.e., with more bins). In line with our earlier synthetic example (Figure 16), this reveals the resolution limit of some records—for example, Piton de la Fournaise appears too coarse under this scheme to support meaningful structure inference.

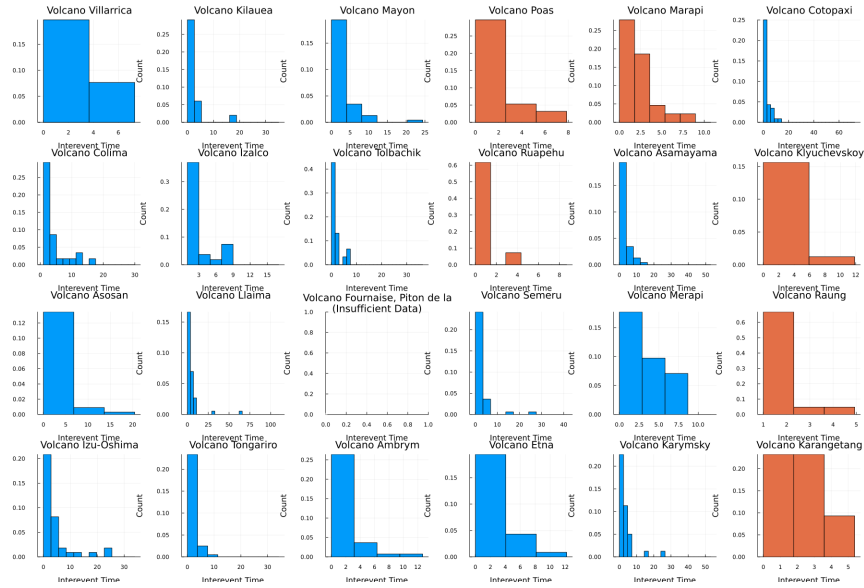


Figure 25. Histograms of inter-arrival times for 24 volcanoes under the maximum-variance condition. Orange subplots indicate insufficient inferential power to reject a randomness null hypothesis. Only records with adjusted bin widths  $\geq$  their Freedman–Diaconis threshold are included.

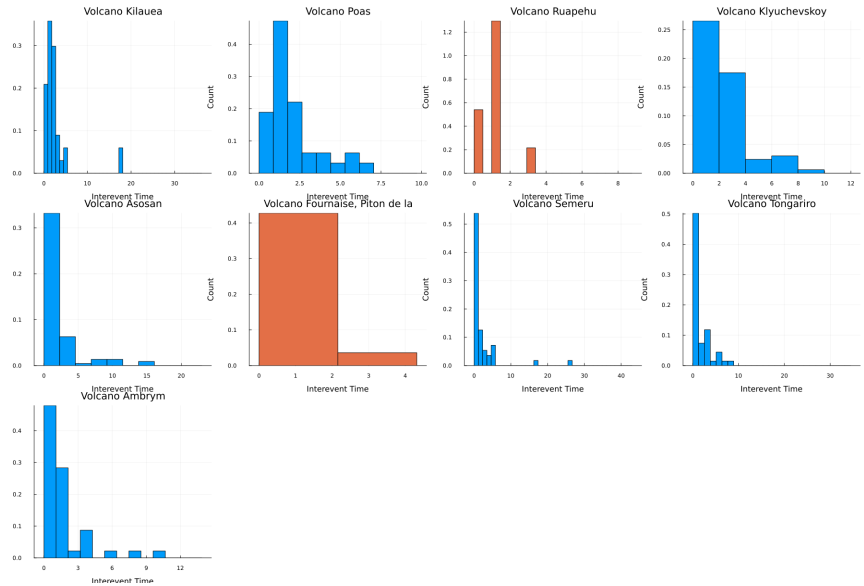


Figure 26. Histograms of inter-arrival times under the median-variance condition, using Asosan as the benchmark. Only records with sufficient Freedman–Diaconis bin width are shown. Orange highlights indicate insufficient inferential power.

Together, these results address a central practical question: given this set of volcanoes, what is the finest temporal resolution at which their eruption structures can be consistently compared? The number of bins in each subplot reflects the expressivity or complexity of our discretized model, how much temporal structure we are attempting to resolve, while the color coding indicates whether sufficient inferential power remains to distinguish that structure from randomness. This framing directly operationalizes the tradeoff introduced in Figure XXX: pushing resolution too far risks trivializing inference, while stopping too early obscures meaningful patterns. Crucially, this balance offers a concrete way to define volcanic completeness: the records that remain unshaded (i.e., not orange) can be considered complete with respect to the level of structural complexity implied by their bin count. This defines a shared, data-informed timescale of comparison, a level of resolution at which inference is both supported by the data and meaningfully interpretable across volcanoes.

While we now know which records can be meaningfully compared, and at what resolution hypothesis testing remains valid, we have not yet demonstrated what formal statistical tests on these structures might look like. This naturally motivates the next stage of analysis, which we address in greater detail in the Discussion. In fact, our final result, which compares volcanoes across different categories of recording quality, also serves as a bridge into that discussion. The interpretation of these cross-category comparisons is not yet straightforward, and the methods required to extend this analysis naturally lead into broader considerations about how to interpret and statistically test the structural comparisons presented above.

### **3.4.5 R5: Comparing records across drastically different qualities of recording.**

While the previous analyses focused on record segments where recording quality was presumed to be comparable, this analysis makes no such assumption. Here, we require only that each record exhibits internal consistency as measured by the Freedman–Diaconis criterion—indicating some stability in correlations and approximate stationarity of the underlying process. Whereas prior comparisons were based on the segments shown in Figure 21, this analysis draws from the longest viable stable portions identified earlier, as illustrated in Figure 20.

Following the procedure outlined in Section 3.3.5, we hold bin density constant across records to ensure a fair basis for comparison. As with variance normalization,

Comparing Across Drastically Different Recording Qualit




Figure 27. Dimensionless intensity comparison across volcanoes with differing record lengths and quality. Each point reflects a stable record segment normalized by bin density and mapped onto the unit interval. Color represents record duration (brighter = longer).

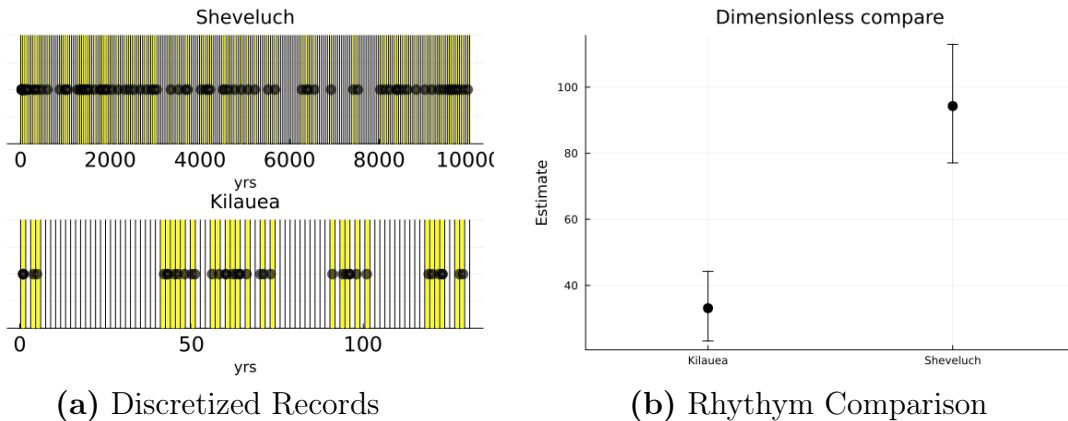
we recheck that each record’s adjusted bin width remains greater than or equal to its Freedman–Diaconis minimum, ensuring that the resulting discretization is not overly sensitive to noise. An additional normalization step rescales each equivalent Poisson process to the unit interval, allowing records of vastly different durations to be compared directly.

In this setting, “intensity” no longer represents the number of eruptions per unit time, but rather a dimensionless measure of burstiness or dispersal. Lower values indicate more clustered or bursty records, while higher values reflect more evenly dispersed eruption timing. In this sense, the comparison is analogous to measuring the rhythmic variability of different systems, a method that could just as easily be applied to heartbeats as to earthquakes. For now we just call the measure rhythm.

The results, shown in Figure 27, reveal a substantial reordering relative to the original intensity estimates in Figure 22. Notably, the point colors reflect record length, with brighter hues representing longer records. This demonstrates that the ranking is not simply a proxy for record duration—precisely the kind of spurious correlation this analysis is designed to avoid.

We can also perform more direct, pairwise comparisons between records. Figure 28(a) shows the discretized eruption records of Kilauea and Sheveluch. Kilauea represents one of the most consistently recorded, frequently active modern volcanoes, while Sheveluch provides a rare, long-lived record despite relatively infrequent eruptions in recent decades. In panel (b), we compare their rhythms,. Despite

the drastically different time spans and event counts, we can say with statistical confidence that Kilauea exhibits a burstier rhythm than Sheveluch. Further insight could be gained by analyzing the structure of their inter-arrival histograms.



*Figure 28.* (a) Discretized records for Sheveluch and Kilauea at resolution which preserves bin-density. Active bins are yellow and the continuous process is shown as markers (b) The rhythm (dimensionless intensity) of Sheveluch and Kilauea.

To further demonstrate the flexibility of this approach, we compare Villarrica, one of the most persistently erupting volcanoes in the modern record, with Mt. Mazama, whose activity spans hundreds of thousands of years and whose last eruption occurred more than 7,000 years ago. Figure 29 shows the application of our stability selector to Mt. Mazama, including the resulting discretized process and inter-arrival time histogram. Given the sparseness of this record, the Freedman–Diaconis bin width is naturally on the order of kiloyears, reflecting the multi-millennial gaps between eruptions. Despite these extreme differences in scale, our cross-scale comparison framework allows for direct alignment between Villarrica and Mt. Mazama. Figure 30 shows their respective discretized records and inter-arrival histograms. When we compare rhythms, we find that Mt. Mazama actually exhibits slightly more bursty behavior than Villarrica, though the difference is less pronounced than in the Sheveluch–Kilauea comparison. These examples illustrate how rhythm allows us to characterize volcanic behavior across dramatically different observational regimes.

Together, these comparisons underscore the versatility of the framework introduced in this chapter. We now turn to the Discussion, where we explore what types of analyses would be most fruitful on these carefully selected and prepared portions of the volcanic record.

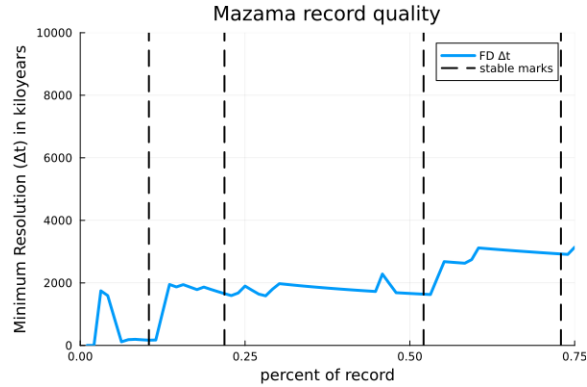
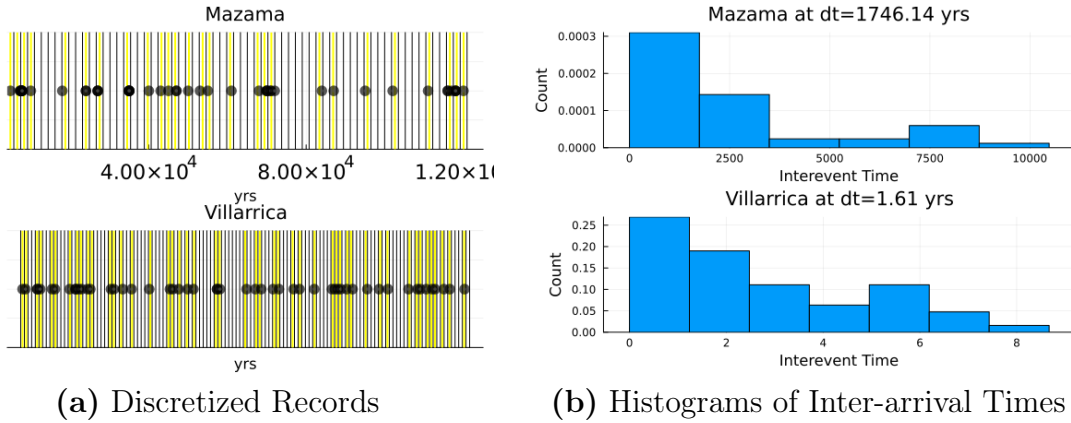


Figure 29. Caption



(a) Discretized Records

(b) Histograms of Inter-arrival Times

Figure 30. (a) Discretized records for Mazama and villarrica at resolution which perserves bin-density. Active bins are yellow and the continuous process is shown as markers (b) Histograms of inter-arrival times with bin-widht set to perserve bin-density.

### 3.5 Discussion

This chapter set out to develop and apply a framework for statistically comparing eruption records, even in the face of gaps, variable quality, and widely differing durations. The results show that, with careful selection and normalization, meaningful inferences can be made even from data that is sparse or inconsistent. By grounding our comparisons in stable portions of each record and aligning analyses via consistent bin widths, densities, and variances, we ensured that the comparisons were fair and interpretable.

Figures 25 and 26 demonstrate that, despite the challenges, many eruption records still contain recoverable structure. However, the amount of usable information is limited. Most records cannot support highly complex models—statistical power is

only maintained for relatively coarse-grained hypotheses. This underscores a key point: when dealing with volcanic data of this kind, formal hypothesis testing must remain general.

In this spirit, the most basic hypothesis we test, "Is the structure of inter-arrival times distinguishable from random?", proves to be both tractable and revealing. For records where the answer is no, we may reasonably interpret the sequence as being effectively random at the tested resolution. For those where the null is rejected, a natural follow-up question arises: why is the pattern non-random? One useful distinction is between sequences whose histograms are monotonically decreasing (consistent with memoryless but bursty behavior) versus those that exhibit distinct modes (suggestive of periodicity or clustering). This simple classification scheme offers a surprisingly insightful window into eruptive dynamics while remaining robust to sparse data.

This classification logic also supports the rhythm-based comparisons presented in the final results. Dimensionless rhythm, derived from normalized Poisson processes, gives us a compact measure of eruptive regularity that is meaningful across vastly different time scales. But more than a summary statistic, rhythm can be paired with structural testing to interpret behavior over a range of temporal aggregations.

In fact, an informative way forward would be to treat the aggregation phase space, that is, the range between a record's minimum resolution (given by Freedman–Diaconis) and its coarsest valid resolution (where the random null can still be rejected at  $\alpha = 0.1$ ), as a diagnostic landscape. By sliding across this range, we can observe how each volcano's structure transitions with increasing resolution. Some may appear random at coarse scales but reveal clustering or periodicity as resolution increases; others may exhibit consistent dispersal patterns. This opens the door to a taxonomy of eruptive behavior, applicable across systems as varied as geysers and flood basalts.

Finally, the hierarchical structure of active bins at different discretization levels hints at a new way to define volcanic episodes. Continuous stretches of active bins across resolutions could be interpreted as nested or cascading eruptive periods, allowing us to formalize the notion of eruptive "events" in a consistent, data-driven way. This may provide a flexible but rigorous language for connecting physical processes to temporal patterns in the record.

### 3.6 Conclusion

This chapter presents a statistical framework for analyzing volcanic eruption records, even when they are incomplete or of varying quality. Our method uses Freedman-Diaconis bin widths to pinpoint reliable data segments, ensuring consistent analysis. By discretizing continuous event times into binary sequences, we effectively manage the bias-variance trade-off and enable robust comparisons of eruption intensity across volcanoes through an equivalent Poisson process. This approach not only standardizes intensity measurements but also quantifies the inferential power available at different resolutions. We demonstrate how maintaining consistent bias or variance across records allows for fair comparisons of both intensity and temporal patterns. Additionally, by normalizing for bin density, our method facilitates meaningful cross-scale comparisons, providing insights into eruptive "rhythm" from records of vastly different durations.

Ultimately, this work provides a robust operational definition of volcanic record completeness, one that is relative to the structural complexity a given dataset can support. It shifts the paradigm from seeking an elusive "perfect" record to strategically leveraging the information that is present. By defining a data-informed timescale of comparison, we provide a means to ask questions commensurate with the available evidence, fostering more rigorous hypothesis testing and a deeper understanding of eruptive dynamics. This framework lays the groundwork for future investigations into the taxonomy of eruptive behavior, offering a flexible yet rigorous language to connect underlying physical processes to the observed temporal patterns in Earth's captivating and hazardous volcanic systems.

## CHAPTER IV DEFORMABLE COLLISION MODEL

### 4.1 Introduction

Granular flows, ubiquitous in natural phenomena from landslides to pyroclastic density currents, present complex rheologies that are inherently scale-dependent. As explored in Chapter 2, understanding the fundamental physical mechanisms governing these systems is crucial, especially when considering the subtle yet significant role of particle deformation in phenomena like thin layer stability. Our previous work demonstrated that even an infinitesimal degree of softness in "rigid" particles allows for serial compression, which profoundly influences contact forces and ultimately impacts macroscopic flow behavior. Specifically, Chapter 2 introduced a novel approach to multi-particle collision modeling by employing the linear dashpot model to constrain the system, thereby enabling unique solutions for previously intractable  $n > 2$  particle elastic deformation problems.

While Chapter 2 focused on leveraging this core concept to explain specific emergent behaviors in monodisperse granular systems, the underlying numerical model developed to achieve this, the Multi-Particle Elastic Collision Model, simply referred to as the deformable collision model, represents a significant independent contribution that warrants a dedicated exploration. Traditionally, modeling elastic deformation in multi-particle systems has been limited by the insufficient constraints provided by conservation laws alone, typically restricting analytical solutions to two-particle interactions. While particle overlap techniques have partially mitigated this by reliably recovering nuanced granular behavior [Garg et al., 2012], they typically yield only averaged effects of the serial compression chains that form during momentum propagation. Critically, these methods cannot resolve the exact contact structure (as illustrated in Figure 33 of this chapter and Figure 8 from Chapter II). This limitation can hinder our understanding of the critical role these compression chains play in setting important, still-phenomenologically understood properties such as the diffusion of granular temperature [Henann and Kamrin, 2014, Kim and Kamrin, 2023], thin layer stability (as shown in Chapter II), or even sound speed propagation in granular material (as hinted at in section 4.5.2). Our approach overcomes this by explicitly modeling deformation through the linear spring model, providing the necessary constraints to resolve unique solutions for momentum and energy across

a variety of one-dimensional particle collision scenarios, even when all material and geometric variables differ.

This chapter is dedicated to detailing the comprehensive development, generalization, and robust implementation of this novel numerical model. We extend the foundational concepts introduced in Chapter 2 by:

- Generalizing the model to accommodate heterogeneous systems, where each of the  $n$  particles can possess unique and variable material properties, including disparate sizes, stiffnesses, and masses.
- Transforming the numerical model into a well-documented and accessible open-source Julia package. This effort underscores the model’s utility as a reusable and extensible software tool for the granular physics community, facilitating further research and application.

The aim of this chapter is to present the theoretical underpinnings, numerical methods, and software architecture of this model, demonstrating its power as a tool for probing the complexities of multi-particle elastic interactions. Through validation of its physical consistency and illustrative applications distinct from the specific scientific question addressed in Chapter 2, we demonstrate its power as a tool for uniquely resolving complex multi-particle elastic interactions and for revealing surprising phenomena in granular dynamics, such as contact ideals and the profound effects of material heterogeneity.

## 4.2 Theoretical Framework

The deformable collision model is based on three foundational assumptions:

1. The particles experience elastic collisions.
2. The elastic forces during collision are modeled through a linear spring model, where the contact force ( $F$ ) is proportional to the inter-particle overlap or compression ( $\Delta x$ ), i.e.,  $F = k\Delta x$ , with  $k$  representing the contact stiffness.
3. Deformations are axisymmetric, meaning that particle boundaries are always equidistant from the center of mass.

In elastic collisions, the final velocities of colliding particles can be determined by conserving total linear momentum and kinetic energy. For a one-dimensional system

of  $n$  particles with masses  $m_i$  and initial and final velocities  $u_i$  and  $v_i$ , respectively, these conservation laws take the form:

$$\text{Momentum: } \sum_{i=1}^n m_i u_i = \sum_{i=1}^n m_i v_i \quad (4.1)$$

$$\text{Energy: } \sum_{i=1}^n \frac{1}{2} m_i u_i^2 = \sum_{i=1}^n \frac{1}{2} m_i v_i^2 \quad (4.2)$$

These are two scalar equations that apply globally to the system, regardless of the number of particles involved. When  $n = 2$ , there are two unknowns ( $v_1, v_2$ ) and two equations ((4.1) and (4.2)), yielding a unique solution. This determinacy vanishes for  $n \geq 3$ . The number of unknowns ( $v_1, \dots, v_n$ ) exceeds the number of constraints. Specifically, there are  $n$  unknowns but still only two equations, making the system underdetermined. As a result, there is a continuum of possible final states consistent with conservation of momentum and energy. Without additional information, such as spatial configuration, timing of pairwise contacts, or assumptions about contact sequences, there is no unique solution.

Luckily, assumptions 2 and 3 allow us to find equations of motion linking the interparticle forces to the distances between them. This is derived in more detail in chapter II, but for two identical particles you arrive at the following equations of motion:

$$a_1 = \frac{k}{m} (2r_0 - (x_2 - x_1)) \quad a_2 = -\frac{k}{m} (2r_0 - (x_2 - x_1)) \quad (4.3)$$

While these equations offer a complete description, they lack analytical solutions and thus require numerical modeling. The central objective of developing our general deformable collision model, the focus of this chapter, involves two key advancements. First, we must generalize the established link between particle position and interparticle force to accommodate particles with variable radii, masses, and stiffnesses. Second, we need a robust method to identify, solely from their center-of-mass positions, which particles within a set form contiguous contacts and thus constitute a 'compression chain' to which these generalized equations apply. We will address these challenges sequentially, beginning with the generalization of the equations of motion to resolve particle boundaries for a given contiguous subset.

### 4.2.1 Generalizing the Equations of motion and resolving particle boundaries

We begin by leveraging a result from Appendix A of Chapter 1 (A.2), which, using our assumption of axisymmetric deformation along the line of centers, expresses all boundaries within a compression chain in terms of its leftmost boundary ( $b_0$ ). Following our convention, for a set of particles  $\{p_1, p_2, \dots, p_n\}$  forming a compression chain, their centers of mass are ordered  $x_1, x_2, \dots, x_n$  from left to right, with corresponding boundaries  $b_0, b_1, \dots, b_n$ . Here  $b_0$  denotes the leftmost boundary of particle 1, and  $b_i$  represents the right boundary of particle  $i$  (which is coincident with the left boundary of particle  $i + 1$  for  $i < n$ ), assuming a known contiguous chain. This leads to the following recurrence relations:

$$\begin{aligned}
 & b_0 \\
 & b_1 = 2x_1 - b_0 \\
 & b_2 = 2x_2 - 2x_1 + b_0 \\
 & b_3 = 2x_3 - 2x_2 + 2x_1 - b_0 \\
 & \vdots
 \end{aligned} \tag{4.4}$$

Observing this pattern, any boundary  $b_k$  in the chain can be generally expressed in terms of the leftmost boundary  $b_0$  and the particle center-of-mass positions  $x_i$  up to  $k$  as:

$$b_k = \sum_{i=1}^k (-1)^{k-i} 2x_i + (-1)^k b_0 \tag{4.5}$$

Next, Newton's Third Law guarantees that the magnitude of the contact force is equal between any two adjacent particles,  $i$  and  $i+1$ . We define partial compression, denoted by  $c_{j,k}$ , as the specific deformation of particle  $j$  resulting from its contact with particle  $k$ . Consequently, the total deformation of an internal particle  $j$ ,  $\Delta x_j$ , is the sum of its partial compressions from its left and right neighbors:  $\Delta x_j = c_{j-1,j} + c_{j,j+1}$ . Under Assumption 2 (the linear spring model), the contact force at an interface  $(j,k)$  is given by  $F_{j,k} = k_j c_{j,k}$ . Newton's Third Law then implies that the stiffness-scaled partial compressions across interfaces are equivalent, meaning  $k_j c_{j,k} = k_k c_{k,l}$  for contiguous contacts along the chain. This fundamental principle establishes how forces propagate and are balanced throughout the compression chain. By systematically applying these force equilibrium relationships at each particle interface, and by leveraging the fact that end member particles (leftmost and rightmost) have only one contact contributing to their total deformation, we can derive a recursive

relation. This allows us to express the total deformation of the first particle,  $\Delta x_i$ , in terms of the total deformations and stiffnesses of subsequent particles, which unfolds through the chain as shown:

$$\begin{aligned}
\Delta x_1 &= \frac{k_2}{k_1} \left( \Delta x_2 - \frac{k_3}{k_2} c_{(2,3)} \right) \\
\Delta x_1 &= \frac{k_2}{k_1} \left( \Delta x_2 - \frac{k_3}{k_2} \left( \Delta x_3 - \frac{k_4}{k_3} c_{(3,4)} \right) \right) \\
&\vdots \\
\Delta x_1 &= \frac{k_2}{k_1} \left( \Delta x_2 - \frac{k_3}{k_2} \left( \Delta x_3 - \frac{k_4}{k_3} \left( \Delta x_4 - \dots \frac{k_{n-1}}{k_{n-2}} (\Delta x_{n-1} - \Delta x_n) \right) \dots \right) \right)
\end{aligned} \tag{4.6}$$

Finally, to link particle position to individual particle deformation, we define the deformation  $\Delta x_i$  for the  $i$ th particle. If  $r_i^0$  represents the relaxed (undeformed) radius of the  $i$ th particle, then the distance from its center of mass  $x_i$  to its leftmost boundary  $b_{i-1}$  would be  $r_i^0$  in an undeformed state. Any deviation implies deformation. Thus, we can express the deformation of particle  $i$  as:

$$\Delta x_i = 2(r_i^0 - (x_i - b_{i-1})), \quad i \in \{1, 2, \dots, n\} \tag{4.7}$$

Putting this all together, we substitute Equation 4.7 for the compressions in Equation 4.6. We then rewrite all  $b_i$ 's in terms of  $b_0$  using the generalized boundary relation (Equation 4.5) and re-arrange to find an equation for the position of the first boundary as a function of particle position, stiffness, and relaxed radii. To simplify the expression, we first define the recursive weight factor:

$$P_i = \prod_{j=1}^i \frac{k_{j+1}}{k_j}, \quad \text{with } P_0 = 1$$

While this simplifies to  $\frac{k_n}{k_1}$ , we leave it in the above form because I feel that better emphasizes the physicality of the situation which is sequential scaling of relative stiffness.

Define the normalization constant:

$$\kappa = \sum_{i=0}^{n-1} P_i$$

Then the leftmost boundary  $b_0$  is:

$$b_0 = \frac{1}{\kappa} \left[ \sum_{i=1}^{n-1} (-1)^{i+1} P_i (r_{i+1} - x_{i+1}) + \sum_{i=1}^{n-1} P_i \left( \sum_{k=1}^i (-1)^{k+1} 2x_k \right) - r_1 + x_1 \right] \tag{4.8}$$

With  $b_0$  determined, we can then use Equation 4.5 to find all of the remaining boundaries within the chain, leveraging our assumption of spherically symmetric deformation. Thus, given an identified compression chain, we have established a method to uniquely determine the full set of internal and external boundaries by applying Assumptions 2 and 3 and invoking Newton’s Third Law. The crucial next step is to identify precisely which subsets of particles form these contiguous compression chains given the broader system’s particle positions.

### 4.2.2 Identifying Chains of Compression

Having established how to resolve particle boundaries within a given compression chain, the next critical challenge is to identify which particles form such chains within a larger system. Given a set of particles  $\{p_1, p_2, \dots, p_n\}$  with their respective positions ( $x_i$ ), stiffnesses ( $k_i$ ), and masses ( $m_i$ ), our goal is to find the partition  $\mathcal{P}$  of these particles that accurately represents all existing compression chains in the system. For instance, in a five-particle system, a partition like  $\mathcal{P} = \{\{1, 2, 3\}, \{4, 5\}\}$  would indicate that particles 1, 2, and 3 are in a contiguous compression chain, followed by a gap, and then particles 4 and 5 form another distinct compression chain. Once this partition is identified, we can precisely determine all relevant particle boundaries using the methods described in the previous section.

**Perhaps one of the most surprising physical insights gleaned from building this model is that the problem of identifying compression chains, and thus accurately modeling the elastic propagation of momentum, ends up being equivalent to an optimization problem.** Specifically, it’s about finding the partition  $\mathcal{P}$  that minimizes the overlap between identified chains. In an ideal, physically consistent scenario, there should be no overlap between distinct compression chains; particles are either part of a chain with defined contacts or separated by a gap. The correct allocation of boundaries, according to our physical constraints of equal inter-particle force and axisymmetric deformation, naturally groups particles together and only allows for gaps where positions explicitly create such deformation-accommodating spaces.

Initially, when solving this problem for simpler systems (e.g., monodisperse particles with identical material properties in chapter II), it seemed possible to enumerate a straightforward set of logical checks to construct the correct particle partition. However, and in the spirit of complete transparency, attempting to build a universally consistent set of rules and conditions for this more general case

(heterogeneous particles with variable properties) quickly led to an unmanageably complex and exhaustive list. Consequently, we adopted a brute-force optimization approach. At each time step, our model calculates every theoretically possible particle partition for the given system. For each potential partition, we then compute the respective boundaries for every chain within that partition using the generalized methods from the previous section. Finally, we select the partition that yields the absolute minimum overlap between chains.

The robustness of this approach is quite remarkable: at each discrete time step, the system’s evolution under elastic deformation is uniquely recovered solely from its current particle positions, material properties, and the previous time step’s velocities. Although a dedicated validation section will provide formal verification, it’s worth noting here that we have analytically validated this logic against two-particle systems. For systems with  $n \geq 3$  particles, where traditional conservation of energy and momentum alone cannot uniquely determine the outcome of elastic collisions, the addition of our deformation constraints fundamentally changes the problem. In our model, contact forces are governed by linear spring laws that relate inter-particle deformation to force. This directly leads to a deterministic system of coupled, continuous differential equations for the particle motions.

Given initial positions and velocities, the system evolves precisely according to Newton’s Second Law, guided by these smooth, linear, and deterministic force-displacement relationships. Because the governing equations are continuous and the force laws are well-defined, standard results from differential equations imply that the solution—that is, the time evolution of particle positions and deformation—is unique and continuous. Physically, this means there is only one way for the system to evolve that is consistent with both the initial conditions and the deformation mechanics. This inherent uniqueness and continuity provides a powerful means of internal validation for our model: for  $n \geq 3$  particles, a critical check on the numerical solution’s consistency is to ensure that the deformation curves evolve continuously over time. Any unphysical ‘jumping’ or discontinuity in these curves would immediately signal that an incorrect partition or boundary allocation was chosen at that specific timestep, indicating a deviation from the system’s unique historical evolution. Thus, our deformable collision model elegantly resolves the inherent ambiguity present in purely conservation-based formulations, ensuring a

physically meaningful and uniquely determined outcome even for complex multi-particle collisions.

### 4.3 Model Implementation and Software

The preceding theoretical framework lays the groundwork for understanding multi-particle elastic collisions by introducing a novel approach to define a deterministic system. Translating this theory into a practical computational tool requires a robust and carefully structured implementation. This section details the architecture, core algorithms, and practical aspects of the numerical model, which has been developed as an open-source Julia package named `DeformableGrains.jl`.

#### 4.3.1 Overall Architecture and Workflow

The `DeformableGrains.jl` package serves as the central module, designed for clarity, shareability, and transparency to facilitate ease of use and reproducibility. It integrates three primary components: `helper.jl`, `collision_model.jl`, and `visualize_results.jl`. This modular design separates concerns, allowing for independent development and easier understanding of different model aspects.

The simulation workflow proceeds through discrete timesteps, orchestrated by the `compression_chain` function located in `collision_model.jl`.

1. **System Setup:** The simulation begins by defining the initial conditions of the particle system. Leveraging Julia’s **multiple dispatch** feature, the model offers two flexible ways to set up the initial state:
  - One method utilizes the `setup_dust_settled` function (in `helper.jl`) to automatically arrange  $N$  particles tangent to one another, with only the leftmost particle receiving an initial velocity. This is particularly useful for controlled experiments focused on analyzing momentum propagation. If gravity is enabled, this setup intelligently pre-positions the particles in a static equilibrium where their compressive forces balance gravitational effects.
  - The second method allows the user to directly specify the initial positions and velocities for all particles, offering maximum flexibility for custom scenarios.

Both setup methods support an optional gravitational field (`g`) and the ability to fix the last particle in space (`stuck`), mimicking a wall.

2. **Time Integration:** The model employs a standard 4th-order Runge-Kutta (RK4) method to numerically integrate the particles' equations of motion over time. This robust numerical scheme is well-suited for systems governed by continuous differential equations, and its reliability has been confirmed through validation against simpler oscillatory systems. The core `F` function serves as the derivative function for the RK4 solver, calculating the accelerations (and thus rates of change of velocities and positions) for all particles at each step.
3. **Iterative Simulation Loop:** The `compression_chain` function drives the simulation through a loop, advancing the system state timestep by timestep. Crucially, within each iteration of this loop, the model dynamically executes the two core theoretical steps detailed in the previous section: identifying compression chains and resolving particle boundaries, which then enable the calculation of inter-particle forces and accelerations.

#### 4.3.2 Core Algorithmic Steps: Implementing the Novelty

The primary innovation of this model lies in its ability to deterministically resolve multi-particle elastic collisions. This is achieved through two tightly integrated algorithmic steps executed at every timestep:

##### 4.3.2.1 Identifying Compression Chains (*find\_chains*)

As discussed, identifying the correct partition of particles into contiguous compression chains is framed as an optimization problem where the objective is to minimize any unphysical overlap between distinct chains. The `find_chains` function in `helper.jl` implements a brute-force approach to solve this, leveraging bitmasking to exhaustively explore all possible partitioning configurations.

The bitmasking strategy operates as follows:

- For a system of  $N$  particles, there are  $N - 1$  possible gaps between adjacent particles.
- Each of these gaps can either be a "cut" (indicating the end of one chain and the start of another) or not.
- This translates to  $2^{N-1}$  total possible configurations. `find_chains` iterates through integers from 0 to  $2^{N-1} - 1$ , where each integer acts as a unique bitmask.

- Each bit within the mask corresponds to a specific gap (e.g., bit 0 for the gap after particle 1, bit 1 for the gap after particle 2, and so on). A '1' at a bit position signifies a cut, while a '0' signifies no cut. This process ensures that every conceivable way to partition the particles is systematically generated and considered.

For clarity, consider an example with  $N = 4$  particles, which presents  $N - 1 = 3$  possible gaps. The following table illustrates all  $2^3 = 8$  bitmask combinations and their corresponding particle partitions:

Table 2. All bitmask combinations and their resulting partitions for  $N = 4$  particles.

bits	Binary	Chains Produced
0	000	[[1,2,3,4]]
1	001	[[1], [2,3,4]]
2	010	[[1,2], [3,4]]
3	011	[[1], [2], [3,4]]
4	100	[[1,2,3], [4]]
5	101	[[1], [2,3], [4]]
6	110	[[1,2], [3], [4]]
7	111	[[1], [2], [3], [4]]

For each candidate partition generated by the bitmasking:

- The `solve_b0` function (described below) is called for each identified chain to determine its leftmost boundary.
- The `satisfies_newton_third_law` function (also described below) performs a crucial physics check to validate the resulting deformed radii for each chain.
- The relative overlap between adjacent chains in the candidate partition is calculated.
- Finally, the partition that yields the absolute minimum inter-chain overlap is selected as the physically correct configuration for that timestep. A small tolerance (`tol`) is used to account for machine precision when determining contact and overlap.

#### 4.3.2.2 *Resolving Boundaries and Calculating Forces* (`solve_b0`, `satisfies_newton_third_law`, `calc_forces`)

Once the optimal compression chain partition is identified, the model proceeds to resolve the precise boundaries of each particle and calculate the forces acting on them:

- `solve_b0`: For each identified compression chain, the `solve_b0` function (in `helper.jl`) uses the generalized theoretical equations (4.5 and 4.8) derived in the previous section to uniquely determine the position of the leftmost boundary ( $b_0$ ). From  $b_0$  and the particle center-of-mass positions, the remaining internal and rightmost boundaries of the chain are then calculated.
- `satisfies_newton_third_law`: This function acts as a critical internal validation step. It verifies that the deformed radii resulting from the resolved boundaries are physically consistent with the model’s assumptions and Newton’s Third Law. Specifically, it checks that particles are not unphysically stretched (deformed radii are not larger than relaxed radii) and that the stiffness-scaled partial compressions (which represent the transmitted force) are equivalent across interfaces within the chain. Any deviation from these conditions signals an unphysical configuration, leading the `find_chains` algorithm to discard that partition.
- `calc_forces`: With the precise particle boundaries and resulting deformations determined, the `calc_forces` function (in `helper.jl`) is responsible for computing the inter-particle forces within each compression chain. It effectively "works backward" from the identified particle deformations (`squish` in the code) and their relative stiffnesses to ascertain the magnitude of the inter-particle forces (`force_L` and `force_R`). These forces are then used, along with the particles’ masses, to calculate the accelerations for each particle according to Newton’s Second Law. While a damping parameter ( $\gamma$ ) is present in the function signature for future extensions, it is not the focus of the current model’s elastic deformation calculations.

Finally, the `F` function (the derivative function for the RK4 solver) orchestrates these steps by calling `find_chains` to obtain the correct partition and then iterating through each identified chain, calling `calc_forces` to sum the accelerations contributed by all active forces in the system, including gravity.

### 4.3.3 Software Tooling and Accessibility

Beyond its core algorithmic implementation, the `DeformableGrains.jl` package also includes utilities for post-processing and visualization. The `visualize_results.jl` file contains functions such as `animate_particles`, which can generate animations and various plots of the simulation data, including particle positions, velocities, and deformations over time. The overarching design principles of clarity, shareability, and transparency, inherent in its Julia package structure, aim to make this novel deformable collision model a valuable and accessible tool for further research in granular physics.

## 4.4 Simple Simulation Walkthrough

To illustrate the practical application of the `DeformableGrains.jl` package, this section provides a step-by-step walkthrough of a basic simulation run. This example, derived directly from `6p_example.jl`, demonstrates how to set up initial conditions, execute the simulation, and visualize key outputs, showcasing the model’s capabilities for analyzing multi-particle elastic collisions.

### 4.4.1 Setting Up the Simulation

A simulation begins by defining the physical parameters of the system. For a simple demonstration, a system of six particles ( $N = 6$ ) is initialized. The leftmost particle is given an initial velocity ( $v = 1$  m/s) towards the right, while the others are initially at rest and tangent to each other. Material properties for each particle, including stiffnesses ( $K$ ), masses ( $M$ ), and relaxed radii ( $R$ ), are defined as vectors, allowing for heterogeneous particle properties as highlighted in the model’s generalization. A total simulation time is also specified.

The setup code, demonstrating these parameter definitions, is as follows:

```
N = 6
v = 1
K = [100,100,1000.0,1000.0,1000.0,100]
M = [1.0,1.0,0.25,0.25,0.25,1.0]
R = [1.0,1.0,.25,.25,.25,1.0]
time = 2
```

### 4.4.2 Running the Collision Model

With the parameters defined, the core simulation is executed by calling the `compression_chain` function. This function encapsulates the entire time-stepping

process, including the RK4 integration, dynamic identification of compression chains, boundary resolution, and force calculations at each timestep. Optional parameters such as the timestep size ( $\Delta t$ ), gravitational acceleration ( $g$ ), and whether the last particle is stuck (`stuck`) can also be set.

The call to run the simulation, capturing its outputs, is:

```
Cs,T,Xs,Vs,Ws =
compression_chain(N,v,K,M,R,time;g=0,dt=0.0001,stuck=false);
```

The function returns several key outputs:

- **Cs**: A collection of the identified compression chains at each timestep.
- **T**: A vector containing all simulated time points.
- **Xs**: A 2D array storing the positions of each particle at every timestep.
- **Vs**: A 2D array storing the velocities of each particle at every timestep.
- **Ws**: A 2D array storing the deformed widths of each particle at every timestep.

#### 4.4.3 Accessing and Visualizing Results

The generated simulation data can be easily accessed for post-processing and analysis. The `6p_example.jl` script demonstrates how to generate various plots to visualize the system’s dynamic behavior.

##### 4.4.3.1 Key Plots

- **Particle Velocities**: A plot of particle velocities over time can be generated to observe momentum transfer and the overall dynamic response of each particle.
- **Particle Deformations**: A plot showing the deformation (initial width minus current width) of each particle over time provides direct insight into how individual particles compress and rebound during collisions. This is a critical output for validating the model’s ability to capture elastic deformation.
- **Compression Chain Length**: The length of the first compression chain through time can be plotted to observe how the overall contact structure evolves, highlighting periods of compression and separation.

These plots utilize Julia’s `Plots` package, directly accessing the `T`, `Xs`, `Vs`, and `Ws` arrays.

#### 4.4.3.2 Animation

For a more intuitive understanding of the particle dynamics, the simulation results can be converted into an animated GIF. The `visualize_results.jl` file provides the `animate_particles` function for this purpose. First, the simulation data is organized into a `DataFrame` and written to a CSV file.

Example of data preparation for animation:

```
H = diff(T)[1] # Timestep size
df = DataFrame(x1=Xs[1,:],...,x6=Xs[6,:], # Positions
              v1=Vs[1,:],...,v6=Vs[6,:], # Velocities
              e1=Ws[1,:],...,e6=Ws[6,:], # Widths/Deformations
              time=T,h=H)
CSV.write("6_particles.csv",df)
```

Then, the `animate_particles` function is called, pointing to the generated CSV file and specifying the output filename for the GIF. This visualization tool allows for the direct observation of particle movements, compressions, and the formation/dissolution of compression chains, providing a qualitative validation of the model's behavior. An example of the dashboard generated by the GIF can be seen in figure 31

Example of animation generation:

```
animate_particles(df,"6p_test.gif",frame_step=100)
```

This walkthrough demonstrates the straightforward process of utilizing the `DeformableGrains.jl` package to simulate and analyze multi-particle elastic collision phenomena, offering a powerful tool for research in granular dynamics.

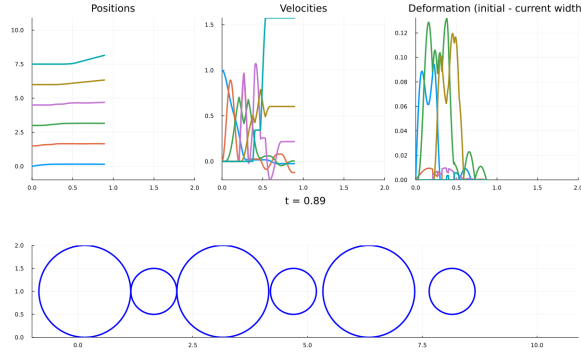


Figure 31. Caption

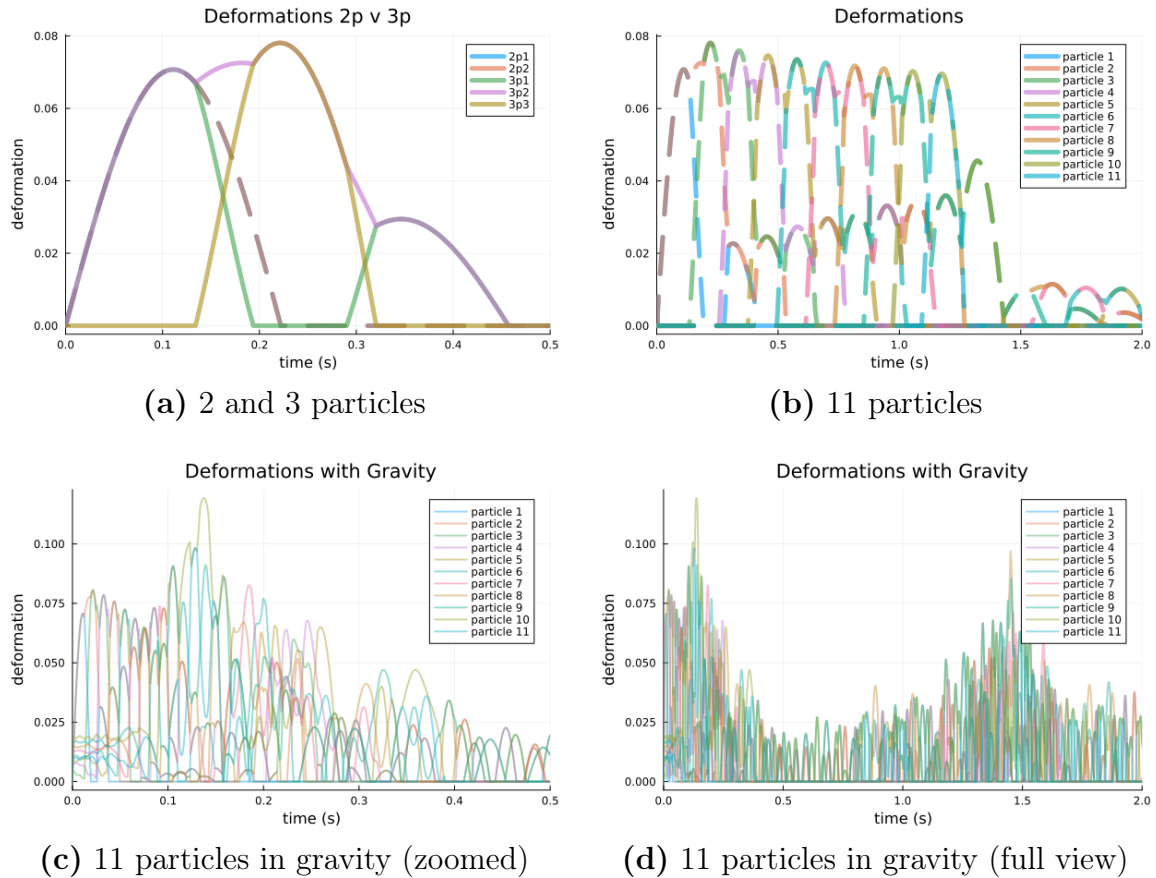
#### 4.5 Validations and Demonstrative Examples

Validation of our numerical model, as mentioned in the theoretical framework, primarily revolves around two critical checks: the conservation of energy and the production of continuous deformation curves. Throughout the simulations presented in this section, changes in kinetic energy remain less than one thousandth of a percent, affirming the model’s adherence to energy conservation. A direct check for the continuity of deformation curves, a key indicator of the model’s internal consistency and uniqueness as derived from differential equation theory, is visually demonstrated in Figure 32.

Figure 32 presents four panels illustrating the model’s performance under various conditions:

- **Panels (a) and (b): Unconfined Systems Without Gravity.** Panel (a) compares the deformation curves of a two-particle system (dashed lines) with a three-particle simulation. A notable immediate change in system behavior is evident: introducing more particles leads to larger total deformation despite identical initial conditions. Panel (b) further demonstrates the model’s consistent and perfect resolution of deformation in a larger system, up to eleven particles.
- **Panels (c) and (d): Extreme Tests with Fixed End and Gravity.** These simulations represent the most challenging tests for our system. Here, the end particle is fixed in place, and gravity is aligned in the direction of motion—analogue to dropping a particle onto a tower of ten other particles. Panel (c) provides a zoomed-in view of the deformation curves,

clearly illustrating their smooth and continuous resolution. Panel (d) shows the system’s larger-scale oscillatory behavior over extended time. These simulations offer an interesting study in complexity: despite being a system of many separable springs, which introduces numerous degrees of freedom, the model robustly recovers an effective larger oscillation that resembles how a single spring with reduced stiffness would respond to forcing under gravity.



*Figure 32.* (a) Deformation curves comparing a two-particle system (dashed lines) with a three-particle system, showing the change in system behavior with increased particle count. (b) Deformation curves for an 11-particle system, demonstrating continuity and resolution across multiple particles. (c) Zoomed-in view of an 11-particle system under gravity, illustrating the fine-grained continuity of deformation curves. (d) Full-time view of an 11-particle system under gravity, revealing a larger-scale oscillatory behavior characteristic of the system’s collective response.

Beyond these fundamental validations, the remainder of this section presents demonstrative examples that leverage the full capabilities of our generalized model.

We first revisit concepts from Chapter 2 concerning monodisperse systems, showing how our more general model can resolve contact structures with greater precision. Subsequently, we present new demonstrations that take full advantage of the model's capability to handle variable mass, stiffness, and size values, illustrating surprising effects arising from heterogeneity.

#### 4.5.1 Contact Ideals

One surprising physical insight gleaned from the development of this model, particularly when applied to simple elastic collisions (uniform size, mass, and stiffness), is the apparent existence of "contact ideals." By this, we mean that for a given system size  $n$ , when properly non-dimensionalized to the timescale of a two-particle collision of equivalent material properties, the evolution of contact structures through time appears identical, regardless of the ratios of mass, stiffness, or initial kinetic energy imbued to the system (excluding pathological cases where energy is zero or deformation exceeds particle widths).

This phenomenon is illustrated in Figure 33, which traces the contact structure experienced by the first particle throughout time. Here, "contact structure" refers to the length of the compression chain that the first particle is currently part of.

These contact ideals hint at a surprisingly conserved behavior behind the complex dynamics governing the propagation of momentum through elastic deformation (which involves translation of kinetic energy to potential energy, formation of gaps, acceleration through those gaps, and momentum transfer dictated by the number of particles currently in contact). This is somewhat surprising given the non-analytic nature of the equations of motion that govern this system. However, in the presence of gravity, these contact ideals vanish, giving rise to much richer behaviors.

Figure 34 illustrates the emergence of a rich behavior in contact structure in the presence of gravity. Two things are particularly noteworthy: (1) The ratio of kinetic energy to system stiffness must be sufficiently high to cause deformation that exceeds what arises in a static spring system at rest under gravity. Otherwise, any perturbations will simply be accommodated by the relaxation of already compressed springs (similar to Figure 32(a) for a purely elastic system). (2) As the energy in the system becomes sufficient to cross gaps, the constant acceleration afforded by the gravitational field introduces enough complexity that ratios of stiffness, mass, and energy once again significantly influence the system's behavior (similar to Figure 32(b) and (c) for higher  $N$  systems).

### 4.5.2 Variable Properties

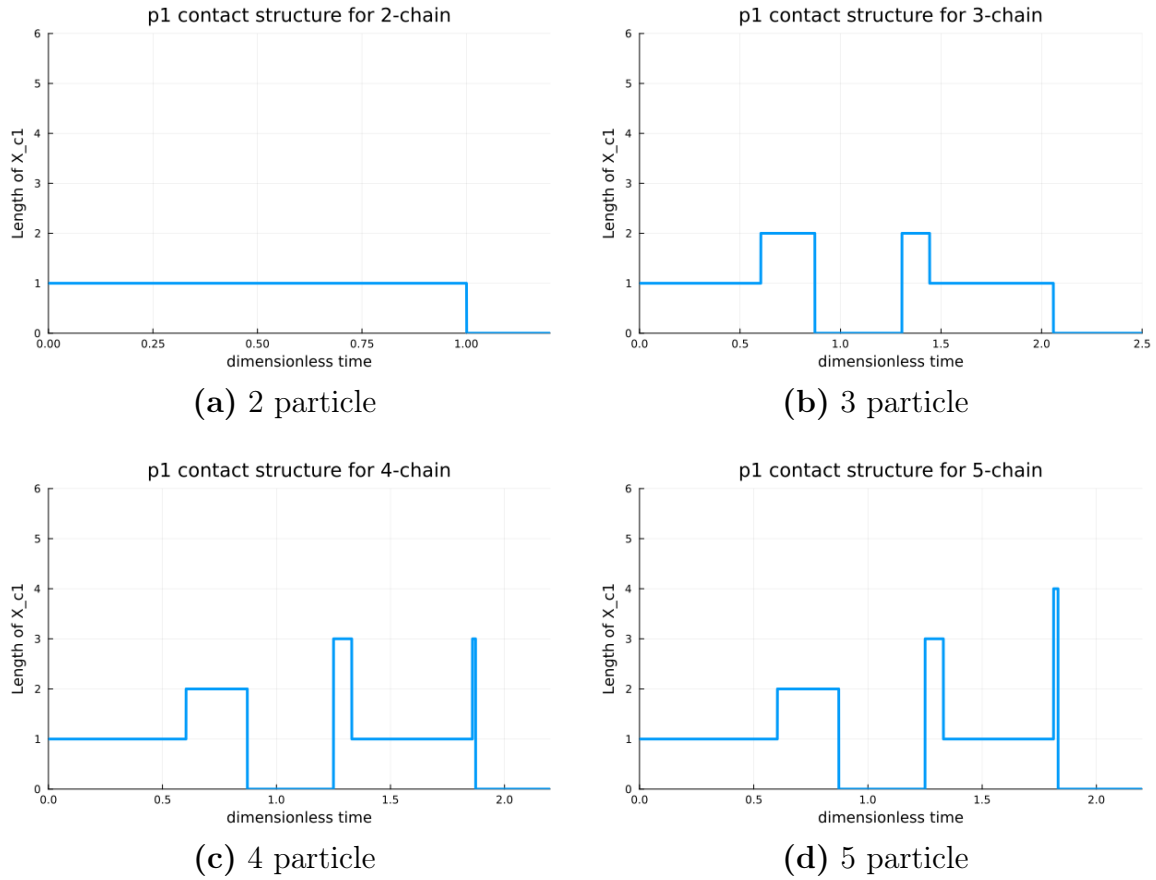
In these demonstrations, we introduce variability in particle stiffness while holding all other material properties constant, and then repeat the analysis for variable mass. Our aim is to investigate (1) how the distribution/ordering of variability affects the system, and (2) how differences in variance impact system behavior. For all these experiments, we use a 6-particle unconstrained system (end particle not fixed in space) with no gravity. The mean system values are held constant across comparisons, and the systems are non-dimensionalized to the timescale of contact for a two-particle system with material properties set to the same mean values. This setup was chosen to explore how variations in material properties can give rise to changes in system behavior even in the simplest settings. While the effects on the system are wide-ranging, we focus on one specific measure: the velocity of the last particle. This allows us to gain insight into how the timescale for propagation of momentum (or sound speed) of the system can change.

**R1) Impact of Ordering:** Figure 35 illustrates the effect of particle ordering. In panel (a), we compare two systems with varying stiffness properties. One system alternates between stiffness values of 50 and 1000 N/m, such that  $K = [50, 1000, 50, 1000, 50, 1000]$ . The other system is "well-sorted" with a stiffness array of  $K = [1000, 1000, 1000, 50, 50, 50]$ . We observe that changing the ordering of stiff particles from alternating to well-sorted generally speeds up the arrival time of momentum to the last particle and ensures a more uniform arrival. In contrast, the alternating scenario ultimately results in less momentum being transferred, as more remains scattered among the trailing particles. Panel (b) shows similar experiments for mass, with lower mass values of 0.25 kg and 1 kg. Here, the effect is swapped: the sorted arrival time is slower, though less dramatically so. However, the difference in the efficiency of momentum transfer is stark. For the sorted particles, much more momentum remains scattered in the trailing particles, while almost all of it is transferred in the alternating scenario. It is notable that the last particle in this setup has a larger initial velocity than the first particle, which is attributed to its smaller mass. A quick check on energy conservation for the variable mass setup confirms that the difference in kinetic energy between these two systems is less than  $-6.49260669849383 \times 10^{-6}$  Joules.


**R2) Impact of Variance:** This next set of experiments keeps the mean material property values the same but compares a system of particles with slightly

varying values to a system featuring a few high-variance particles. For example, for stiffness, we compare  $K1 = [1000, 800, 1000, 800, 1000, 800]$  (representing relatively uniformly stiff particles) with  $K2 = [1300, 1300, 100, 1300, 100, 1300]$  (a system with some higher stiffness particles mixed with two very "squishy" particles). A similar check is performed for mass with  $M1 = [1, 0.9, 1, 0.9, 1, 0.9]$  and  $M2 = [1.3, 1.3, 0.25, 1.3, 0.25, 1.3]$ . The results are shown in Figure 36.

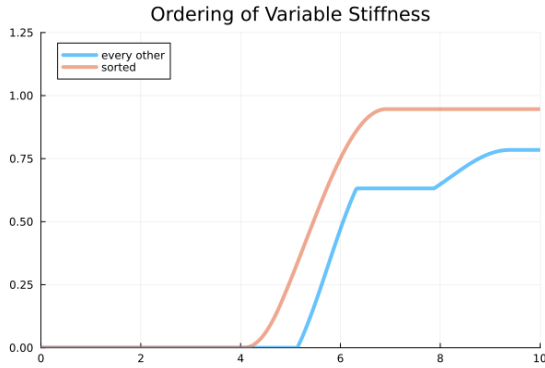
We observe a dramatic change in the momentum propagation time between the low and high variance stiffness cases. This hints at the potential for a few anomalously stiff materials mixed into a well-known granular material to cause significant deviations from expected behavior. We see less of an effect on variable mass, but high variance in mass does affect the uniformity of momentum transfer, which is distributed over many subsequent collisions rather than quickly propagating.




*Figure 33.* These subplots illustrate the contact structure for particle one, showing the length of the compression chain in contact with particle one through time. These systems are unconstrained (last particle not fixed in space) and not subject to gravity. Time is non-dimensionalized to the timescale of contact for a two-particle system of equivalent properties. (a) The 2-particle system sets the non-dimensionalization of time. (b) In the 3-particle system, at most two other particles can be in contact with particle one. (c) 4-particle system. (d) For systems with 5 or more particles, this represents the limiting contact structure from the first particle’s point of view, with no further changes observed as the number of particles continues to increase.



*Figure 34.* Illustration of the complex contact structure behavior that arises in the presence of gravity. The y-axis represents contact length, and the x-axis represents non-dimensionalized time.




(a) Stiffness Ordering




(b) Mass Ordering

Figure 35. Impact of particle ordering on the final particle’s velocity. (a) Compares alternating and well-sorted stiffness distributions. (b) Compares alternating and well-sorted mass distributions.



(a) Stiffness Variance



(b) Mass Variance

Figure 36. Impact of material property variance on the final particle’s velocity. (a) Compares low and high variance stiffness distributions. (b) Compares low and high variance mass distributions.

## 4.6 Discussion

The comprehensive validation presented in the preceding section confirms the robust performance of our deformable collision model. The demonstrated adherence to energy conservation and, critically, the continuous evolution of deformation curves (Figure 32) affirm the model's internal consistency and its capability to deterministically resolve multi-particle elastic collisions. This continuity provides crucial validation for systems with three or more particles, where traditional approaches often lead to ambiguous solutions.

While the model exhibits excellent and expected physical behavior for systems up to 11 particles, the brute-force chain identification technique used in the current implementation imposes an  $O(2^n)$  computational complexity. This exponential scaling limits local simulations to approximately 15 particles before encountering significant performance constraints. However, this practical limitation is somewhat mitigated by existing literature suggesting that a particle length scale of around 10 particles frequently defines the mesoscale behavior in many granular materials [Zhang and Kamrin, 2017]. Despite this contextual relevance, enhancing the computational efficiency of the chain identification algorithm beyond brute force remains a critical avenue for future work to expand the model's applicability to larger systems.

Looking ahead, the current framework readily accommodates the introduction of viscous damping. Implementing this involves using the change in boundary widths from previous timesteps to derive the speed of compression, which then informs the damping force. This extension is crucial for exploring energy dissipation in granular systems, bridging elastic and inelastic models. Furthermore, it will stabilize simulations under gravity, where the current absence of dissipation causes noisy and sensitive long-term behavior, thus enabling more accurate studies of sustained granular dynamics under constant external forces.

Beyond validation, the demonstrative examples revealed compelling insights into granular dynamics. The emergence of "contact ideals" in homogeneous systems, where non-dimensionalized contact structures evolve identically regardless of initial kinetic energy or specific material property ratios (Figure 33), highlights a surprising underlying regularity in an otherwise non-analytic system. More significantly, the investigations into heterogeneous systems uncovered surprisingly large effects on momentum propagation stemming from the ordering and variance of particle properties, even when average system values were held constant. As illustrated in

Figures 35 and 36, subtle changes in particle arrangement or the introduction of a few highly variant particles can dramatically alter the speed and uniformity of momentum transfer. While these initial demonstrations involved simplified systems and did not constitute an exhaustive parameter sweep, they are sufficient to show a profound sensitivity of system-level behavior to the micro-scale distribution of material properties.

In essence, the model's validation confirms its accuracy and consistency, while the demonstrative examples reveal critical, often counter-intuitive, dynamics that arise from elastic deformation and material heterogeneity in multi-particle systems.

#### **4.7 Conclusion**

This chapter developed a novel numerical model to deterministically simulate elastic collisions involving three or more particles, a problem traditionally left ambiguous by conservation laws alone. Our approach, rooted in modeling particle deformation through a linear spring mechanism, allowed for the unique resolution of these complex multi-particle interactions.

Key insights derived from the model include the existence of "contact ideals" in homogeneous systems and a surprising sensitivity of momentum propagation to the ordering and variance of material properties in heterogeneous particle assemblies. Validated for accuracy and consistency, this model is encapsulated in an accessible open-source Julia package, designed for clarity, shareability, and transparency, thereby serving as a valuable tool for investigating the intricate dynamics of granular materials.

## APPENDIX

### SUPPORTING INFORMATION FOR CHAPTER 2: THE ROLE OF COMPRESSIONAL DYNAMICS IN SETTING THE SCALE-DEPENDENT RHEOLOGY OF GRANULAR FLOWS: AN EXPLANATION FOR THE EMERGENCE OF THIN LAYER STABILITY

#### A.1 Resolving Boundaries from Center of Mass

Recovering the boundaries for a one-dimensional system of deformable particles based only on center of masses can be framed by two questions; Can we identify which particles are compressed in chains (i.e. where are the gaps), and then, within each chain of compressed particles, can we recover deformed boundaries? We start with the second question. Throughout this section assume that numerical and alphabetical indexing order a particles position from left to right. Thus particle  $i$  or 1 is left-adjacent to particle  $j$  or 2.

##### A.1.1 Recovering boundaries within a compressed chain

If we have a compressed chain of  $n$  particles and the positions of their center of masses  $\{x_1, x_2, \dots, x_n\}$ , then there exist  $n + 1$  boundaries  $\{b_0, b_1, \dots, b_n\}$ , and without loss of generality we make  $b_0$  the left most boundary in the chain. These boundaries are recoverable with two constraints: (1) axisymmetry of deformation and (2) Newton's third law. By axisymmetric deformation we can write a recurrent relationship for the particle boundaries.

$$\begin{aligned} & b_0 \\ b_1 &= x_1 + (x_1 - b_0) \\ b_2 &= x_2 + (x_2 - b_1) \\ b_3 &= x_3 + (x_3 - b_2) \\ & \vdots \\ b_n &= x_n + (x_n - b_{n-1}) \end{aligned} \tag{A.1}$$

Simply put, the above relation states that the right boundary of the first particle must be the same distance from the center of mass as the left boundary, and that this right boundary is the left boundary of the next particle, and so forth...(remember we start by assuming these particles are in contact and part of a compressive chain). This relationship allows for every boundary to be written as the particle's center of

masses and the left most boundary.

$$\begin{aligned}
& b_0 \\
& b_1 = 2x_1 - b_0 \\
& b_2 = 2x_2 - 2x_1 + b_0 \\
& b_3 = 2x_3 - 2x_2 + 2x_1 - b_0 \\
& \vdots \\
& b_n = \sum_{i=1}^n (-1)^{n-i} 2x_i + (-1)^n b_0
\end{aligned} \tag{A.2}$$

Since the total compression of the chain is known from information about the center of mass positions  $(x_n - x_1)$  so too is the length of the chain, which provides us enough information to find all the boundaries for odd chains of compression. We provide an example with a three particle chain, by subtracting  $b_0$  from both sides of the equation above:

$$\underbrace{b_3 - b_0}_{\text{length}} = 2x_3 - 2x_2 + 2x_1 - 2b_0$$

All of the positions are known, therefore we can solve for  $b_0$  and use the recurrence relationship to find all particle boundaries. Interestingly, this is not the case for even chains of compression because the  $b_0$  term on the right hand side cancels out. This surprising result makes more sense when considering a chain of two particles, which can have identical centers of mass positions and chain length but different boundaries by swapping deformation of each particle; odd particles break this symmetry. However, Newton's third law removes this degree of freedom (in a two particle chain the deformation of the two particles must be identical).

We incorporate Newton's third law by linking the distance between the the position and boundaries to the compression in our linear dashpot model. The total compression of partilce  $i$  is given by the difference between its relaxed length ( $2r_0$ ) and its deformed length  $((b_i - x_i) + (x_i - b_{i-1}))$ . Once again using the assumption of axisymmetric deformation this can be simplified into equation A.3

$$\Delta x_i = 2 \cdot (r_0 - (x_i - b_{i-1})) \tag{A.3}$$

The total compression of some particle  $i$  will have two contributing factors, a partial compression from its left and right neighbors, represented as  $c_{(i-1,i)}$  and  $c_{(i,i+1)}$  respectively (equation A.4)

$$\Delta x_i = c_{(i-1,i)} + c_{(i,i+1)} \tag{A.4}$$

The first and last particle of the chain have the property that their total compression is equivalent to their partial compression:

$$\begin{aligned}\Delta x_1 &= c_{(1,2)} \\ \Delta x_n &= c_{(n-1,n)}\end{aligned}\tag{A.5}$$

It is worth noting that the axisymmetric assumption for compression guarantees that  $b_i - x_i = x_i - b_{i-1}$  not that  $c_{(i-1,i)} = c_{(i,i+1)}$ . In other words, while the total deformation is symmetric the contributions from the left and right neighboring particles to this total deformation do not have to be equal. We can now construct a similar recurrence relationship to equation A.1, but using compression instead. By Newton's third law, the deformation of the first particle must be equal to the deformation of the second particle minus the partial compress caused by the third particle. However the partial compression of the second and third particle must be equal to the total compression of the third particle minus its partial compression from the fourth, and so on...

$$\begin{aligned}\Delta x_1 &= \Delta x_2 - c_{(2,3)} \\ \Delta x_1 &= \Delta x_2 - (\Delta x_3 - c_{(3,4)}) \\ \Delta x_1 &= \Delta x_2 - (\Delta x_3 - (\Delta x_4 - c_{(4,5)})) \\ &\vdots \\ \Delta x_1 &= \Delta x_2 - (\Delta x_3 - (\Delta x_4 - \dots (\Delta x_{n-1} - \Delta x_n) \dots))\end{aligned}\tag{A.6}$$

The last partial compressive term can just be represented as a total compression by equation A.5 . Using equation A.3 we can rewrite all of the  $\Delta x_i$ 's in terms of boundaries, and then using equation A.1, we can rewrite all of these boundaries in terms of  $b_0$  and solve! An example is given below for a chain of four particles:

$$\begin{aligned}\Delta x_1 &= \Delta x_2 - \Delta x_3 + \Delta x_4 \\ 2(r_0 - (x_1 - b_0)) &= 2((r_0 - (x_2 - b_1)) - \dots \\ &\quad (r_0 - (x_3 - b_2)) + (r_0 - (x_4 - b_3))) \\ r_0 - x_1 + b_0 &= r_0 - x_2 + 2x_1 - 2b_0 - r_0 + x_3 - \dots \\ (2x_2 - 2x_1 + b_0) &+ r_0 - x_4 + 2x_3 - 2x_2 + 2x_1 - b_0 \\ b_0 &= \frac{1}{4}(-x_4 + 3x_3 - 5x_2 + 7x_1)\end{aligned}$$

From this boundary we can find all of the others and we have achieved the desired result of recovering all deformations from just center of mass. The pattern for the

general solution is already clear in the solution for just 4 particles (alternating sum of odd coefficients). In general the solution for  $b_0$  is given by equation A.7

$$b_0 = \frac{1}{n} \left( (-1)^n r_0 + \sum_{i=1}^n \left( (-1)^{i+1} (2(n-i) + 1) x_i \right) \right) \quad (\text{A.7})$$

We can now recover the boundaries from chains of compressed particles, what remains is to identify the chains.

### A.1.2 Identifying chains of compressed particles from center of mass positions

As mentioned above, once we know which sets of particles are in contact, we can uniquely identify particle boundaries, however identifying which sets of particles form a compressive chain from only center of mass data is quite challenging. There are many different ways a system of particles could be grouped into sets of compressed chains which satisfy both the assumptions of axisymmetry and Newton's third law. In order to uniquely identify the chains of compressed particles from center of mass position information alone, we add an additional constraint shown in equation A.8 .

$$c_{(i,j)} \leq 2r_0 - \Delta x_i \quad (\text{A.8})$$

One way to interpret this inequality is that given the positions of particle  $i$  and  $j$ , the boundary of particle  $j$  will be as compressed as possible if particle  $i$  has no further particles to its left. Thus, it is not necessarily a new constraint so much as a useful inequality which stems directly from newtons third law. It is useful because it allows us to find a threshold for "closeness" beyond which we can guarantee particle contact. If we know particle  $k$  is in contact with some particle  $l$ , we can make the same assumption about particle  $k$ 's partial compression and if the distance between particle  $j$  and  $k$ 's center of mass is smaller than the some of their maximal compressive distance, particle  $j$  can be guaranteed to be a part of particle  $k$ 's compressive chain. In this way the chains are built in an order of "these particles at least must be in contact". Particles that are part of a chain might be missed in a first sweep but they will be picked up in a later sweep because when encountered again, more information is known about the contacts of the particles on the other side (these would have been established in a previous sweep). All that remains is to find a starting point. Given a set of particle positions and our two assumptions which two particles , if any, are in contact.

Not surprisingly the particles whose center of masses are closest are guaranteed to be in contact. The conceptual argument is fairly simple; in order for two particles to be sufficiently deformed that they are no longer in contact, their positions to their neighboring particles must be even closer. However for the sake of robustness we present a proof by contradiction using our favorite inequality (equation A.8).

Assume to the contrary that particles  $x_i$  and  $x_j$  have the minimum distance between their centers, yet are not in contact. Thus, the distance between particles  $x_i$  and  $x_j$  must be greater than the sum of their radii:

$$x_j - x_i > \frac{2r_0 - \Delta x_i}{2} + \frac{2r_0 - \Delta x_j}{2} \quad (\text{A.9})$$

The deformations of particle  $x_i$  and  $x_j$  do not have to be equal, however we can add the minimum of the two twice, and maintain the inequality. Without loss of generality, assume that  $\Delta x_i = \min\{\Delta x_i, \Delta x_j\}$ . Thus equation A.9 simplifies to:

$$x_j - x_i > 2r_0 - \Delta x_i \quad (\text{A.10})$$

Now we can use our important inequality. The deformation of particle  $x_i$  is maximized (boundary is minimized) when  $x_h$  is not compressed by any particle to its left. By assumption, particles  $x_i$  and  $x_j$  are not in contact and therefore inequality can be preserved and the deformation can be written in the simple two particle case, that is we can write  $\Delta x_i$  as  $(2r_0 - (x_i - x_h))$ :

$$x_j - x_i > 2r_0 - (2r_0 - (x_i - x_h)) = x_i - x_h \quad (\text{A.11})$$

However this means that the distance between particles  $x_i$  and  $x_j$  are not the minimum and we have our contradiction.

## A.2 Picking a time Scale for Impulse Calculations


Implicit in our results is the definition of a timescale over which we calculate our impulses for each simulation. This question of choosing a timescale is important for describing the different phases of interactions which arise in deformable particle collisions. For our simple one dimensional collision model, we begin by creating two aspirational distinctions, mobile grain-neighbor interactions and mobile grain-chain interactions. Mobile grain-neighbor interactions take place at a timescale which captures the initial interactions between the mobile grain and the first particle in the chain. These dynamics will still be influenced by contacts beyond one particle

lengthscale, but the focus is on this initial exchange of momentum. However, when the contacts extend beyond one particle lengthscale (when we have a series of deformed particles) some of the mobile grain's energy will be transferred down the chain. Not only does this reduce the energy in the mobile grain's immediate neighbor interaction, it also creates dynamics which will be resolved with the mobile grain at a later time. These dynamics include gravitational acceleration into gaps created by the elastic propagation of deformation, and the rebounding of the initial momentum. The dynamics which resolve on this timescale are referred to as mobile grain-chain interactions.

In the two particle simulation, mobile grain-neighbor and mobile grain-chain interactions are synonymous. In the simulations with longer chains (greater than ten particles) these interactions are well separated. However, in the transition these interactions overlap, forming a constructive interference which makes identifying a consistent and physically intuitive timescale for the characteristic impulse difficult. We attempt to illustrate the phenomenon in Figure A.37. In the two particle simulation, there is no separation between the mobile grain-neighbor and chain interactions. In the twenty particle simulation, we see the initial interactions well isolated, and the interactions which occur up to  $t/t_0=1.5$  seem to be the limiting expression of the mobile grain neighbor-interactions in long compression chains. However, in the four particle simulation, we see an example of these mobile grain-neighbor interactions (third hump) occurring on timescales which are likely still influencing the particle-neighbor interactions. The three particle case is the most extreme example, with the mobile grain-chain interaction completely overlapping with the second packet of the mobile grain-particle envelope.

While it is tempting to use these contact structures to inform temporal bounds for impulse calculation, it is difficult to tie these to physical intuition and make them consistent for the two particle case, the transition lengths, and the long particle chains. For example, using the first dimensionless time greater than one where the contact structures swing from the total chain length to one, produces a remarkable fit for the transition zone (Figure A.38), however the physical intuition is murky, and it is difficult to reconcile with original two particle scenarios and the longer particle chains.

A simpler choice for impulse timescale would be the periodicity for the mobile grain to collide and rebound. However, the idea of rebound becomes complicated



*Figure A.37.* A clear separation of mobile grain-particle and mobile grain-chain interactions develops as chain length increases. It arises similar to a separation of constructive interference.

once the rebound effects become partitioned into increasingly separated mobile grain-particle and mobile grain-chain interactions. However, what remains clear is the initial compression, so we identify the time, within  $3t_0$  at which the chain is maximally compressed, which represents half of the periodicity and then double this value. The bound of  $3t_0$  is placed to ensure we are looking at compressions directly tied to the initial collision and not effects from compounding compression and gravity later in the simulation. Figure A.39 shows an example of this time being selected for a three particle simulation, and these are the temporal bounds used to produce our impulse calculations in figure 10. The fact that this approach does not directly reconcile with the contact structures, and the fact that the three particle simulation is a unique overlap of different contact dynamics, is our best guess as to why it fails to capture the three particle scenario. We believe the fact that these tricky transition lengths are of lengthscales typically associated with the granular mesoscale is of no coincidence. Important future work would combine information encoded in the contact structures and physical intuition to find a consistent definition of the impulse timescale which also sheds light on the evolution of dynamics within these transition lengths.

### A.3 The role of the restitution coefficient

When doing the DEM validation in the background section (figure 4 we set the restitution coefficient to zero as a validation technique to isolate the role of system geometry and ensure that the global analytical descriptions in Equations 5 and 6 can

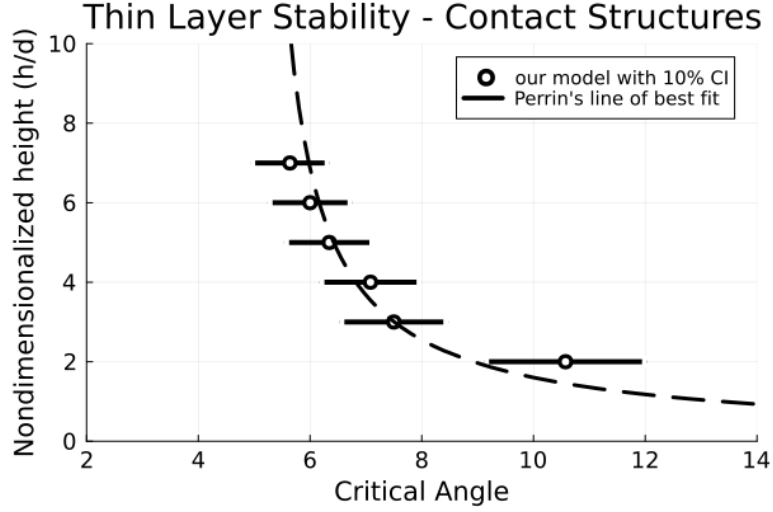


Figure A.38. Critical angles for simulations of two to seven particles when using contact structures to pick temporal bounds for impulse calculation.

accurately replicate the results obtained from the DEM approach. However, in our simulations producing the reported results, we use a restitution coefficient of 0.21, consistent with Quartier’s experimental measurements for his ‘rigid’ cylinders. His work shows that the value appears to be primarily determined by the shock and did not vary noticeably between felt, rubber-coated, or brass rollers (Quartier 2000).

It is important to address the role  $e_n$  can play in setting the scaling. Observing the impact of  $e_n$  on the final result is unfortunately not analytically tractable because it requires isolating  $e_n$  in the following system of equations:

$$\mathcal{E} = \left| \frac{v_0 - v_s}{v_0} \right|^2 = \left( 1 - \sqrt{\left( \frac{1}{2} \right)^2 + \left( e_n \frac{\sqrt{3}}{2} \right)^2} \right)^2$$

$$\frac{\mathcal{E}}{1 - \mathcal{E}} \geq \frac{2 \sin^2\left(\frac{\pi}{12} - \frac{\theta}{4}\right)}{\sin(\theta)}$$

However, we can still capture the effect by conducting a numerical analysis of the sensitivity to  $e_n$  through multiple system runs to produce the following result. This shows that  $e_n$  does not set the scaling behavior but largely produces a horizontal translation:




Figure A.39. Visualization of contact structure (blue), forces felt (orange), and compression of the mobile grain for a 4 particle chain. The purple dashed line indicates where the particle is maximally compressed, which sets the timescale for our impulse calculation.

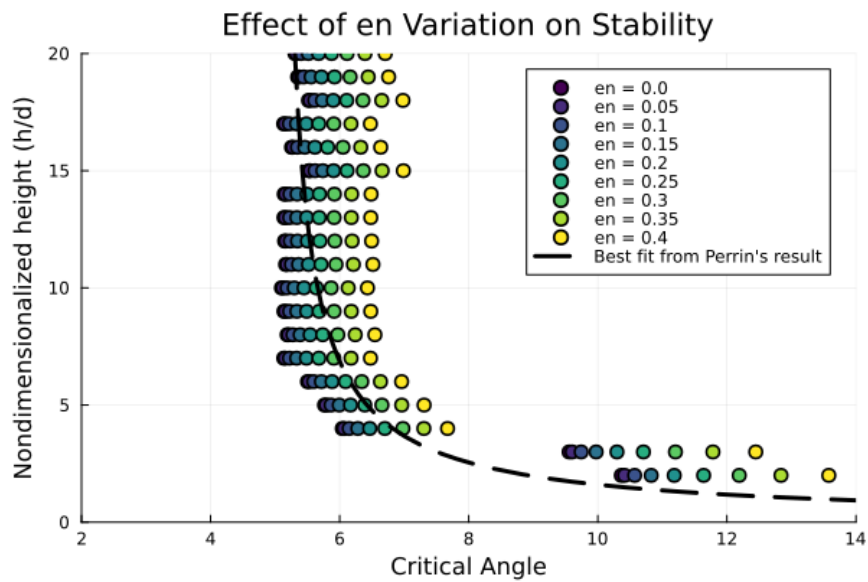


Figure A.40. Sensitivity of our model to different values for the restitution coefficient

## Bibliography

- L. Amarsid, J. Y. Delenne, P. Mutabaruka, Y. Monerie, F. Perales, and F. Radjai. Viscoinertial regime of immersed granular flows. *Physical Review E*, 96, 2017.
- B. J. Andrews, E. Venzke, E. Cottrell, S. Sennert, K. Bennis, and E. Crafford. The volcanoes of the world database. *Bulletin of Volcanology*, 2025. doi: 10.5479/si.GVP.VOTW5-2025.5.3. URL <https://volcano.si.edu/>. Manuscript Number: BUVO-D-24-00143R1, For submission to BV special issue on databases in volcanology.
- P. Bak, C. Tang, and K. Wiesenfeld. Self-organized criticality. *Phys. Rev. A*, 38: 364–374, Jul 1988. doi: 10.1103/PhysRevA.38.364. URL <https://link.aps.org/doi/10.1103/PhysRevA.38.364>.
- M. S. Bebbington and W. Marzocchi. Stochastic models for earthquake triggering of volcanic eruptions. *Journal of Geophysical Research: Solid Earth*, 116(B5), 2011. doi: <https://doi.org/10.1029/2010JB008114>.
- L. Bocquet, A. Colin, and A. Ajdari. Kinetic theory of plastic flow in soft glassy materials. *Physical Review Letters*, 103, 2009.
- F. Boyer, E. Guazzelli, and O. Pouliquen. Unifying suspension and granular rheology. *Physical Review Letters*, 107:1–5, 2011.
- S. G. Coles and R. S. J. Sparks. Extreme value methods for modelling historical series of large volcanic magnitudes. In H. M. Mader, S. G. Coles, C. B. Connor, and L. J. Connor, editors, *Statistics in Volcanology*, volume 1 of *Special Publications of IAVCEI*, pages 47–56. Geological Society, London, 2006. doi: 10.1144/IAVCEI001.4. URL <https://sp.lyellcollection.org/content/specpubgsl/1/1/47.full.pdf>. ISBN: 978-1-86239-208-3.
- D. R. Cox and V. Isham. *Point Processes*. Chapman & Hall, London, 1980. URL <https://www.scirp.org/reference/referencespapers?referenceid=667960>. Cited as a foundational reference for Poisson processes and their applications.
- F. Da Cruz, S. Emam, M. Prochnow, J. N. Roux, and F. Chevoir. Rheophysics of dense granular materials: Discretesimulation of plane shear flows. *Physical Review E*, 72, 2005.

- J. Ewart. An assessment of volcanic threat and monitoring capabilities in the united states: Framework for a national volcano early warning system. Open-File Report 2005-1164, U.S. Geological Survey, 2005. URL <https://pubs.usgs.gov/of/2005/1164/>.
- Y. Forterre and O. Pouliquen. Flows of dense granular media. *Annual Review of Fluid Mechanics*, 40:1–24, 2008.
- D. Freedman and P. Diaconis. On the histogram as a density estimator:  $L^2$  theory. *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete*, 57(4):453–476, 1981. doi: 10.1007/BF01025868. URL <https://link.springer.com/article/10.1007/BF01025868>.
- R. Garg, J. Galvin, T. Li, and S. Pannala. *Documentation of open-source MFI<sub>X</sub>–DEM software for gas-solids flows*. 2012.
- Global Volcanism Program. Volcanoes of the world (v. 5.3.0). Online database, 2025. URL <https://volcano.si.edu/>. Distributed by Smithsonian Institution. Updated 17 July 2025.
- P. Guttorp and M. L. Thompson. Estimating second-order parameters of volcanicity from historical data. *Journal of the American Statistical Association*, 86(415):578–583, 1991.
- D. L. Henann and K. Kamrin. Continuum thermomechanics of the nonlocal granular rheology. *International Journal of Plasticity*, 60:145–162, 2014.
- H. M. Jaeger, S. R. Nagel, and R. P. Behringer. Granular solids, liquids, and gases. *Reviews of modern physics*, 68(4):1259, 1996.
- D. Jessop, K. Kelfoun, P. Labazuy, A. Mangeney, O. Roche, J.-L. Tillier, M. Trouillet, and G. Thibault. Lidar derived morphology of the 1993 lascar pyroclastic flow deposits, and implication for flow dynamics and rheology. *Journal of Volcanology and Geothermal Research*, 245:81–97, 2012.
- C. Johnson, B. Kokelaar, R. M. Iverson, M. Logan, R. LaHusen, and J. Gray. Grain-size segregation and levee formation in geophysical mass flows. *Journal of Geophysical Research: Earth Surface*, 117(F1), 2012.

- K. Kamrin, K. M. Hill, D. I. Goldman, and J. E. Andrade. Advances in modeling dense granular media. *Annual Review of Fluid Mechanics*, 56:215–240, 2024.
- S. Kim and K. Kamrin. A second-order non-local model for granular flows. *Frontiers in Physics*, Volume 11 - 2023, 2023. ISSN 2296-424X. doi: 10.3389/fphy.2023.1092233. URL <https://www.frontiersin.org/journals/physics/articles/10.3389/fphy.2023.1092233>.
- K. Kiyosugi, C. Connor, S. Sparks, S. Croweller, S. Brown, L. Siebert, T. Wang, and S. Takarada. How many explosive eruptions are missing from the geologic record? analysis of the quaternary record of large magnitude explosive eruptions in japan. *Journal of Applied Volcanology*, 4, 07 2015. doi: 10.1186/s13617-015-0035-9.
- B. Kokelaar, R. Graham, J. Gray, and J. W. Vallance. Fine-grained linings of leveed channels facilitate runout of granular flows. *Earth and Planetary Science Letters*, 385:172–180, 2014.
- A. Mangeney, F. Bouchut, N. Thomas, J.-P. Vilotte, and M.-O. Bristeau. Numerical modeling of self-channeling granular flows and of their levee-channel deposits. *Journal of Geophysical Research: Earth Surface*, 112(F2), 2007.
- W. Marzocchi and L. Zaccarelli. A quantitative model for the time-size distribution of eruptions. *Journal of Geophysical Research: Solid Earth*, 111(B4), 2006. doi: <https://doi.org/10.1029/2005JB003709>.
- L. Papadopoulos, M. A. Porter, K. E. Daniels, and D. S. Bassett. Network analysis of particles and grains. *Journal of Complex Networks*, 6:485–565, 2018. ISSN 2051-1329.
- P. Papale. Global time-size distribution of volcanic eruptions on earth. *Scientific Reports*, 8:6838, 2018. doi: 10.1038/s41598-018-25286-y. URL <https://www.nature.com/articles/s41598-018-25286-y>. Open Access, CC BY 4.0.
- P. Papale, W. Marzocchi, and D. Garg. Global volume distribution for subaerial volcanism on earth. *Journal of Geophysical Research: Solid Earth*, 126(6): e2021JB021763, 2021.

- H. Perrin, M. Wyart, B. Metzger, and Y. Forterre. Nonlocal effects reflect the jamming criticality in frictionless granular flows down inclines. *Physical Review Letters*, 126, 2021.
- O. Pouliquen. Scaling laws in granular flows down rough inclined planes. *Physics of Fluids*, 1999. doi: 10.1063/1.1416884.
- O. Pouliquen and N. Renault. Onset of granular flows on an inclined rough surface: Dilatancy effects. <http://dx.doi.org/10.1051/jp2:1996220>, 6, 06 1996. doi: 10.1051/jp2:1996220.
- L. Quartier, B. Andreotti, S. Douady, and A. Daerr. Dynamics of a grain on a sandpile model. *Physical Review E*, 62(6):8299, 2000.
- O. Reynolds. Xxix. an experimental investigation of the circumstances which determine whether the motion of water shall be direct or sinuous, and of the law of resistance in parallel channels. *Philosophical Transactions of the Royal Society of London*, 174:935–982, 1883. doi: 10.1098/rstl.1883.0029. URL <https://royalsocietypublishing.org/doi/abs/10.1098/rstl.1883.0029>.
- F. M. Rocha, C. G. Johnson, and J. M. N. T. Gray. Self-channelisation and levee formation in monodisperse granular flows. *Journal of Fluid Mechanics*, 876:591–641, 2019.
- J. Rougier, R. S. J. Sparks, K. V. Cashman, and S. K. Brown. The global magnitude-frequency relationship for large explosive volcanic eruptions. *Earth and Planetary Science Letters*, 482:621–629, Jan. 2018. doi: 10.1016/j.epsl.2017.11.015.
- P. M. Sadler. Sediment accumulation rates and the completeness of stratigraphic sections. *The Journal of Geology*, 89(5):569–584, 1981. doi: 10.1086/628623. URL <https://doi.org/10.1086/628623>.
- A. Saichev and D. Sornette. “universal” distribution of interearthquake times explained. *Phys. Rev. Lett.*, 97:078501, Aug 2006. doi: 10.1103/PhysRevLett.97.078501. URL <https://link.aps.org/doi/10.1103/PhysRevLett.97.078501>.
- D. W. Scott, R. A. Tapia, and J. R. Thompson. Nonparametric Probability Density Estimation by Discrete Maximum Penalized- Likelihood Criteria. *The Annals of*

*Statistics*, 8(4):820 – 832, 1980. doi: 10.1214/aos/1176345074. URL <https://doi.org/10.1214/aos/1176345074>.

R. Sparks, H. Pinkerton, and G. Hulme. Classification and formation of lava levees on mount etna, sicily. *Geology*, 4(5):269–271, 1976.

T. Wang, M. Schofield, M. Bebbington, and K. Kiyosugi. Bayesian modelling of marked point processes with incomplete records: Volcanic eruptions. *Journal of the Royal Statistical Society Series C: Applied Statistics*, 69(1):109–130, 10 2019.

S. Wiemer and M. Wyss. Minimum magnitude of completeness in earthquake catalogs: Examples from alaska, the western united states, and japan. *Bulletin of the Seismological Society of America*, 90(4):859–869, 08 2000.

L. Zhang, Y. Wang, and J. Zhang. Force-chain distributions in granular systems. *Phys. Rev. E*, 89, 2014.

Q. Zhang and K. Kamrin. Microscopic description of the granular fluidity field in nonlocal flow modeling. *Physical Review Letters*, 2017.