

This is a repository copy of *Robustness of steady state and stochastic cyclicity in generalized coalescence-fragmentation models*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/208716/>

Version: Published Version

Article:

Fagan, Brennen orcid.org/0000-0002-8451-920X, MacKay, Niall orcid.org/0000-0003-3279-4717 and Wood, A. Jamie orcid.org/0000-0002-6119-852X (2024) Robustness of steady state and stochastic cyclicity in generalized coalescence-fragmentation models. European physical journal b. 21. ISSN 1434-6028

<https://doi.org/10.1140/epjb/s10051-024-00654-y>

Reuse

This article is distributed under the terms of the Creative Commons Attribution (CC BY) licence. This licence allows you to distribute, remix, tweak, and build upon the work, even commercially, as long as you credit the authors for the original work. More information and the full terms of the licence here:

<https://creativecommons.org/licenses/>

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.



Robustness of steady state and stochastic cyclicity in generalized coalescence-fragmentation models

Brennen T. Fagan^{1,2,a} , Niall J. MacKay^{1,b}, and A. Jamie Wood^{1,3,c}

¹ Department of Mathematics, University of York, Heslington Lane, York YO10 5DD, UK

² Leverhulme Centre for Anthropocene Biodiversity, University of York, Heslington Lane, York YO10 5DD, UK

³ Department of Biology, University of York, Heslington Lane, York YO10 5DD, UK

Received 13 October 2023 / Accepted 2 February 2024
© The Author(s) 2024

Abstract. Processes of coalescence and fragmentation are used to understand the time-evolution of the mass distribution of various systems and may result in a steady state or in stable deterministic or stochastic cycles. Motivated by applications in insurgency warfare we investigate coalescence-fragmentation systems. We begin with a simple model of size-biased coalescence accompanied by shattering into monomers. Depending on the parameters this model has an approximately power-law-distributed steady state or stochastic cycles of alternating gelation and shattering. We conduct stochastic simulations of this model and its generalizations to include different kernel types, accretion and erosion, and various distributions of non-shattering fragmentation. Our central aim is to explore the robustness of the steady state and gel-shatter stochastic cycles to these variations. We show that an approximate power-law steady state persists with the addition of accretion and erosion, and with partial rather than total shattering. However, broader distributions of fragment sizes typically vitiate both the power law steady state and gel-shatter cyclicity. This work clarifies features shown in coalescence/fragmentation model simulations and elucidates the relationship between the microscopic dynamics and observed phenomena in this widely applicable interdisciplinary model type.

1 Introduction

One of the most prominent quantitative results in political science is due to the physicist Lewis Fry Richardson [1], who observed that deaths caused by other humans are well-approximated by power laws across a wide range of sizes, with an exponent of around 2.5 for murders (1 death) up to small wars (1000 deaths), and a smaller exponent around 1.5 for larger wars. A body of recent work, exemplified by [2] for the FARC insurgency in Colombia, and summarized in [3], has demonstrated that similar distributions occur for event-size distributions for a wide range of modern wars and insurgencies, and the different exponents have been reconciled in [4]. Although Richardson's Law has been challenged – there are cases where the drop-off is faster than expected [5] – it remains a broadly applicable truth.

Why is this so? The belief is that event sizes approximately follow the distributions of the sizes of the human (typically insurgent) groupings which cause them. Distributions which are approximately power-law across most of their range are common in a wide range of phe-

nomena and can result from a range of generative processes including, most prominently, preferential attachment [6]. More broadly, processes of coalescence, or (in finite populations) of coalescence balanced by fragmentation (CF), produce distributions that have approximately power-law steady states. Most simply, a multiplicative kernel, in which the probability of two groups coalescing is proportional to the product of their sizes, yields a power-law exponent of 2.5 (see below). Yet it would seem rather too strong to believe that the complexity of human grouping for deadly conflict is captured by so simple a model. More reasonable, *a priori*, is some form of universality class – that a wide range of such models, with varied rules for coalescence and fragmentation, all produce steady-state distributions with exponents around the observed values. Furthermore we may conjecture that the common structure of observed distributions results from a broad sense in which human conflict, especially at the small scales below those of nation-states, is self-organizing. There is a wide literature exploring this using deterministic mean-field (DMF) methods, summarized in [7] which presented analysis of various generalizations of (1) below, and to which we refer the reader for additional background.

From the perspective of applications to complex systems, the crucial requirement is to understand what

^a e-mail: brennen.fagan@york.ac.uk (corresponding author)

^b e-mail: niall.mackay@york.ac.uk

^c e-mail: jamie.wood@york.ac.uk

features are necessary for genuinely emergent phenomena. Beyond the steady state, the most salient observed features are *gelation*, the absorption of most of the population into a single large group, and the *shattering* of large groups into individuals. Again this can be studied mathematically in the DMF approach [8]. However, we have already reported that in stochastic simulations of a fixed population a new phenomenon occurs, that of *stochastic gel-shatter cycles* [9], in which steady gelation is followed by stochastic shattering, in certain suitable parameter regimes of models with multiplicative size-biased kernels [9]. We term these ‘stochastic cycles’ as although they do not have determinate cycles they are *on average* repeating, and on a characteristic timescale – unlike the chaotic behaviour of a strange attractor, for example. Stochastic cyclic phenomena have been previously reported by McKane and Newman [10].

The purpose of the present paper is to perform stochastic simulations of a variety of CF processes. Whilst our primary concern is the phenomena – the robustness of an approximately power-law steady state and the occurrence and extent of stochastic cyclicity – we also need to understand how these are observed through the statistical methods used to impute power laws to data. Thus we are concerned not with analytical results, which typically require some degree of idealization, but with what reported data, via its analysis, are telling us about the nature of the stochastic generating processes.

We begin with size-biased coalescence accompanied by shattering fragmentation, and then extend to include: different types of kernel; accretion and erosion (‘Becker–Döring dynamics’); and a range of non-shattering fragmentation processes. Whilst these processes are motivated by microscopic rules with a clear justification in our primary application area of operations research, our work is much more generally applicable either in principle or via direct analogy to alternate microscopic behaviour.

Alongside reporting our own simulations we draw together a range of contributory results from the literature. CF models have been used in a wide variety of systems, ranging from the physical interactions between asteroids and dust [11, 12] to probabilistic [13], economic [14], biological [15] and social structures such as the insurgent warfare with which we began [2, 16]. Previous expectations have typically been for fragmentation to balance coalescence in such models, yielding a steady state that follows a truncated power-law distribution (*i.e.* a power-law distribution with an exponential cutoff). Recent literature has questioned this assumption. With coalescence, resupply, and sedimentation without fragmentation, oscillations in the DMF kinetics have been observed [17]. Hopf bifurcations have been observed both in [17] and when employing Becker–Döring mechanics [18]. Deterministic oscillations also occur in systems with fragmentation and when using kernels of the form $K(i, j) = (i/j)^a + (j/i)^a$ [19].

However, and crucially for the present paper, a DMF treatment is not enough. In fully stochastic simulations we see steady states becoming fragile, with some

time-asymmetric stochasticity (and thus moving away from detailed balance), and in some regimes the new phenomenon of stochastic gel-shatter cyclicity, which must now be included alongside deterministic cycles and power-law steady states as a known stable outcome of a CF process. Thus it is only with simulations of the full stochastic process, with statistical identification both of the steady-state distribution and of stochastic cyclicity, that we can describe the full phenomenology of these models.

The classical approach begins with a DMF treatment of coalescence-fragmentation systems of the form

$$\frac{dn_k}{dt} = \frac{1}{2} \sum_{i=1}^{k-1} K(i, (k-i))n_i n_{k-i} - n_k \sum_{i=1}^{\infty} K(i, k)n_i + \frac{1}{k} \sum_{i=k+1}^{\infty} F(i, k)n_i - n_k \sum_{i=1}^k F(k, i), \quad (1)$$

where n_k is the density of clusters of size $k \geq 1$, $K(i, j)$ is the (symmetric) coalescence rate kernel for clusters of sizes i and j to form clusters of size $i + j$, and $F(i, k)$ is the fragmentation rate kernel from clusters of size i to clusters of size k . These systems are known to yield (truncated) power-law distributions when in steady state, $\frac{dn_k}{dt} = 0 \forall k$ [16, 20, 21]. We adopt as our basic model the multiplicative coalescence and shattering fragmentation kernels

$$K_{\text{mult}}(i, j) = \hat{K} \frac{ij}{M^2}, \quad (2)$$

and

$$F_{\text{mult}}(i, k) = \hat{F} \frac{i}{M} \delta_{k,1}, \quad (3)$$

where \hat{K} and \hat{F} are the constant reaction rates, normalized for system size M , the maximum possible number of monomers in the closed system. Normalization by M reflects a probabilistic interpretation: whenever a cluster needs to be chosen, choose a constituent monomer from all possible monomers in the system. The cluster that monomer is currently a part of (including but not limited to clusters of size 1) is the cluster chosen. We will also use $N = \sum_i n_i$ to measure the number of groups (aggregated clusters) in the system. When this system has a steady state, its power law distribution (probability density function) can be found analytically, with an exponent of $\alpha = 2.5$.

Central to our approach, in the light of the results on oscillatory behaviour noted above, is empirically to identify cyclic departures from steady state, so we introduce a cyclicity order parameter \mathcal{K} , defined by

$$\mathcal{K} = \frac{\sum_t \text{sgn}(k_{\text{max}}(t) - k_{\text{max}}(t-1))}{t}, \quad (4)$$

where sgn is the sign function and k_{max} is the size of the largest cluster in the system. This lies between -1

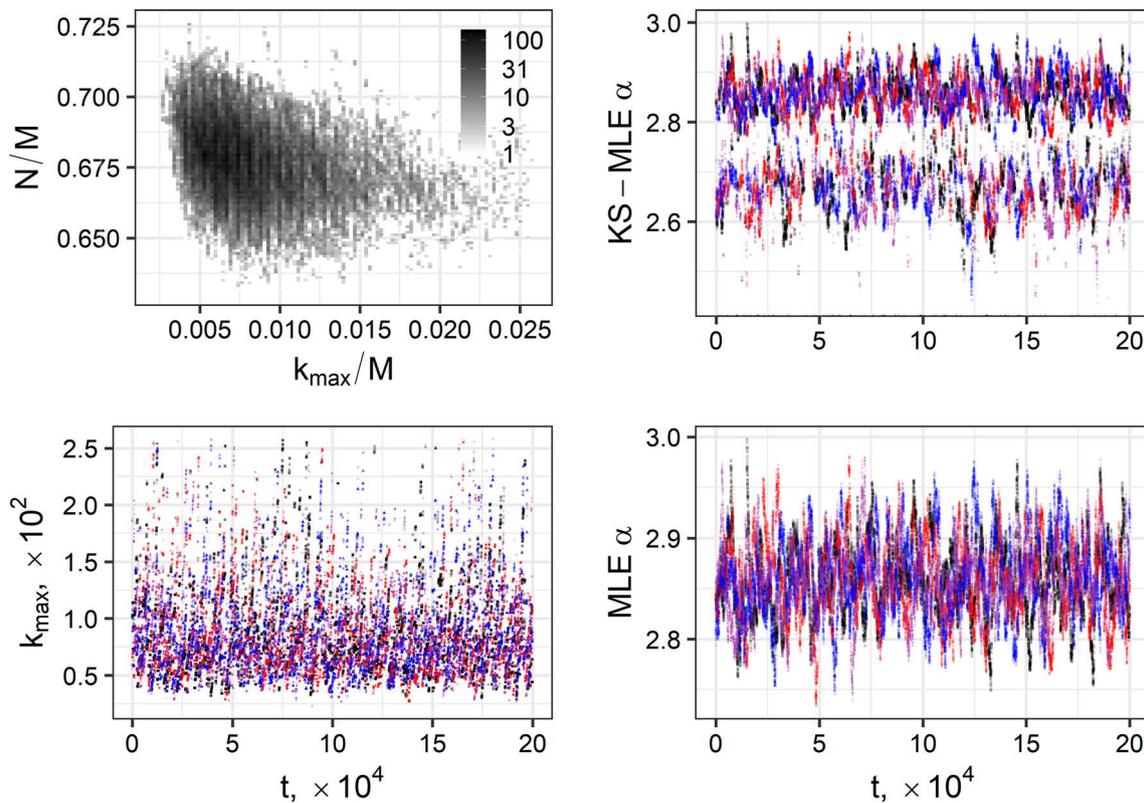


Fig. 1 Coalescence and fragmentation summary statistics, $M = 10^4$, $\hat{F} = 0.20$. Four simulations were conducted (black, blue, red, and purple). From top-left and proceeding counter-clockwise, we present a heatmap of the locations of the simulations in the $(\frac{k_{max}}{M}, \frac{N}{M})$ plane, the time series of k_{max} , the time series of the MLE α estimate, and the time series of the KS-MLE α estimate. We have trimmed large (> 5) and small (< 1) estimates of the KS-MLE α , corresponding to failures to fit the data

and $+1$, and is the proportion of computational time steps for which the largest cluster becomes larger, minus that for which it becomes smaller. In the steady state, or in regimes of symmetric stochastic variation, \mathcal{K} is approximately zero. In regimes of slow build-up followed by sudden fragmentation, by contrast, \mathcal{K} becomes significantly different from zero, and in extreme cases approaches $+1$. Negative \mathcal{K} are seldom observed and never persistent. Thus \mathcal{K} measures the amount of asymmetric stochastic cyclicity in the system, and thereby diagnoses slow-fast cycles, analogous to some relaxation oscillations such as the ‘Oregonator’ [22]. The empirical \mathcal{K} is naturally accompanied by the theoretical dimensionless number r , the ratio of the characteristic time-to-gelation to the characteristic time-to-fragmentation [9].

This paper is structured as follows. We begin with a discussion of our summary statistics in Sect. 2, needed in order to frame and analyse the problem. With these well understood, in Sect. 3 we study variations in the construction of the model including different kernel types and a wide range of parameters, enabling us to establish our base model. In Sect. 4 we study what happens when we add various secondary processes to a particular case of the multiplicative-kernel base model that has a clear steady state with low cyclicity, test-

ing the robustness of the system to changes in the nature of the underlying dynamics. In Sect. 5 we add the same secondary processes to a particular case of the multiplicative-kernel base model that has distinctive gel-shatter cyclicity, again to explore its robustness. We conclude with a discussion of the implications in Sect. 6, with a particular emphasis on problems in operations research.

2 Statistical methodology

To study the behaviour of the system as it evolves over time, we need to first establish a standard method of how we will statistically analyse data emerging from our microscopic simulations. The method we adopt is to fit a power-law distribution at each time step to not be averaging raw data. We choose power-law distributions (rather than more complicated distributions) to study the robustness of existing results in the literature (but see the appendix) [e.g. 2, 23]. We then need to understand how our power-law distribution summary statistics behave.

To fit the power-law distribution we use maximum likelihood estimation for the estimation of the expo-

Table 1 Summary of robustness results in steady-state coalescence and fragmentation. In all cases, $M = 10^4$, $\text{Pr}(\text{Frag.}) = 0.20$, and four simulations were used. The ranges given for KS-MLE and MLE estimates of power-law α contain 95% of empirical results. Perturbations used are explained in the text. We exclude large (> 5) and small (< 1) estimates of the KS-MLE α from our intervals

| Perturbation | Cyclicality \mathcal{K} | KS-MLE ($x_{\min} > 1$) α | MLE α |
|--|---------------------------|------------------------------------|----------------|
| Steady-state | 0.123 | [2.568, 2.754] | [2.788, 2.932] |
| Accretion (3, Unique) | 0.560 | [2.246, 2.685] | [1.907, 2.135] |
| Erosion (3, Unique) | 0.009 | [2.849, 3.047] | [3.122, 3.229] |
| Accretion and Erosion (3, Unique) | 0.266 | [2.341, 2.769] | [2.105, 2.256] |
| Accretion (3) | 0.564 | [2.248, 2.682] | [1.906, 2.136] |
| Erosion (3) | -0.011 | [2.118, 4.848] | [3.528, 3.635] |
| Accretion and Erosion (3) | 0.179 | [2.468, 2.806] | [2.178, 2.315] |
| Coalescence Stick-breaking | 0.094 | [2.534, 2.768] | [2.725, 2.829] |
| Fragmentation Stick-breaking | 0.463 | [1.890, 2.166] | [2.241, 2.402] |
| Coalescence and Fragmentation Stick-breaking | 0.239 | [2.208, 2.393] | [2.298, 2.369] |
| CRP $\theta = 1.20$ | 0.454 | [2.075, 2.257] | [2.274, 2.376] |
| CRP $\theta = 1.50$ | 0.247 | [2.299, 2.426] | [2.434, 2.511] |
| CRP $\theta = 1.80$ | 0.157 | [2.459, 2.620] | [2.640, 2.734] |

ment and the maximum. The maximum likelihood estimator for the exponent (sometimes called the Hill estimator when working with continuous data) is asymptotically normal and consistent [24, 25], and it has seen widespread use due to seminal work [24] and widespread implementation [e.g. 26, 27]. For the minimum, we either use the smallest possible value $x_{\min} = 1$, the monomer of the system, (method ‘MLE’) or Kolmogorov–Smirnov maximum likelihood estimation (‘KS-MLE’) [24], which selects the minimum that corresponds to the power-law distribution with the least Kolmogorov–Smirnov distance (the maximum difference between the cumulative distribution functions) to the empirical data. In effect, KS-MLE fits only the power-law tail of the data, while MLE fits a power-law to the entirety of the data. (We also examined the performance of power-law distributions with exponential cutoffs and observed very little difference in fitted summary statistics.)

To characterize the behaviour of the estimators, we examine how they work in a model system with $M = 10^4$ and $\hat{F} = 0.2 = 1 - \hat{K}$. (Note that these give $r = 2.5 \times 10^3$; steady states are typically observed if $r > 10^3$ [9].) Four simulations of this system can be seen in Fig. 1, whose four panels show: a heat map of the locations the system visits (after burn-in) in $(k_{\max}/M, N/M)$ space; k_{\max} over time; the KS-MLE α values over time; and the MLE α values over time. In practice, this system has a small amount of cyclicality $\mathcal{K} = 0.123$ on average across the simulations, which can be seen upon close examination of the time series. The region explored is quite narrow, with k_{\max}/M within [0.0023, 0.0554] and N/M within [0.6318, 0.7261]. Of these simulations, 99% of time steps sampled remained within [0.0035, 0.0276] and [0.6427, 0.7106] respectively, consistent with steady-state and noise behaviour (in comparison to cyclic behaviour). The average cyclicality and 95% intervals for KS-MLE and MLE α are collected in the first row of Table 1, which collects analogous val-

ues for each model we later consider in Sect. 4. (We take large (> 5) and small (< 1) estimates of the KS-MLE α to indicate failures to fit the data.)

Despite the proximity to the steady state, the system has a distinctive two-layer pattern in the fitted KS-MLE α which is simply explained. Recall that the MLE x_{\min} is always set to 1, but the KS-MLE x_{\min} is allowed to vary to optimize the fit, and the KS-MLE α is highly sensitive to it. The higher values arise when $x_{\min} = 1$ is preferred by the KS-MLE algorithm, while lower values arise when $x_{\min} \geq 2$. Looking at the ‘violin’ distribution plots in Fig. 2, we first note that where MLE and KS-MLE both use $x_{\min} = 1$, the MLE α is equal to the KS-MLE α ; this occurs in around 62% of samples. Where the MLE and KS-MLE x_{\min} differ, there is no obvious relationship between the MLE α and KS-MLE α , although there is a correlation (of 0.454). Henceforth, when we refer to KS-MLE α , we mean cases with $x_{\min} > 1$.

The bias of the fitted exponents away from the expected 2.5 is not due to finite size, see Fig. 3. Instead, it is intrinsic to the system. To check this, we performed a simple experiment. We took the (truncated, normalized) theoretical steady state of the corresponding CF system and used it as a probability distribution to create 1000 random partitions of M , each constructed by drawing a sequence of clusters, with the last cluster truncated when the total population reached M . Results are shown in Fig. 4 and agree well with Fig. 2.

This point is of real importance: there is a substantial difference between the power-law exponent obtained in the calculation of the theoretical solution and the power-law exponent obtained by the simulation, due to the finite nature of the system that the solution is expected to describe. The exponents obtained through the MLE and KS-MLE power-law fits indicate that the latter has a positive bias away from the theoretical solution’s exponent, and studies that rely on these methods of fitting should expect this positive bias to be present.

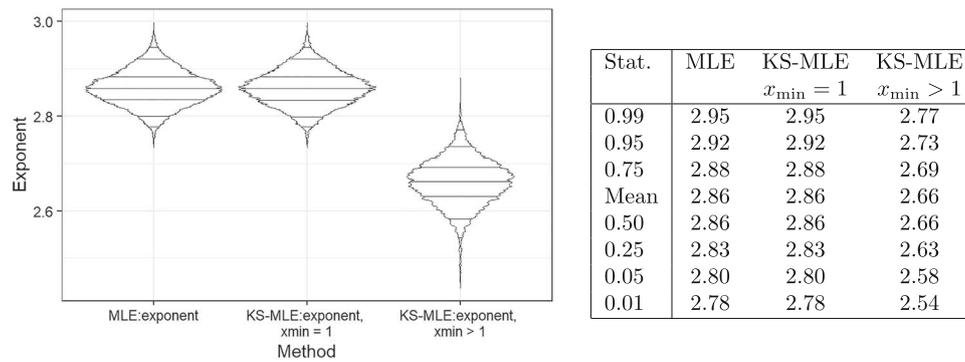


Fig. 2 Coalescence and fragmentation exponent summary statistics. We present the distribution of estimated exponents for the 80,000 samples plotted in Fig. 1 as violin plots for the fitted exponents according to each method (left) with the quantiles (horizontal lines on violin plots) and means (right). Violin plots here are mirrored (probability) density plots plotted with 1/8th the usual bandwidth to preserve the detailed edges [28]. MLE α and KS-MLE α agree when the KS-MLE $x_{\min} = 1$. The theoretical result indicates $\alpha = 2.5$

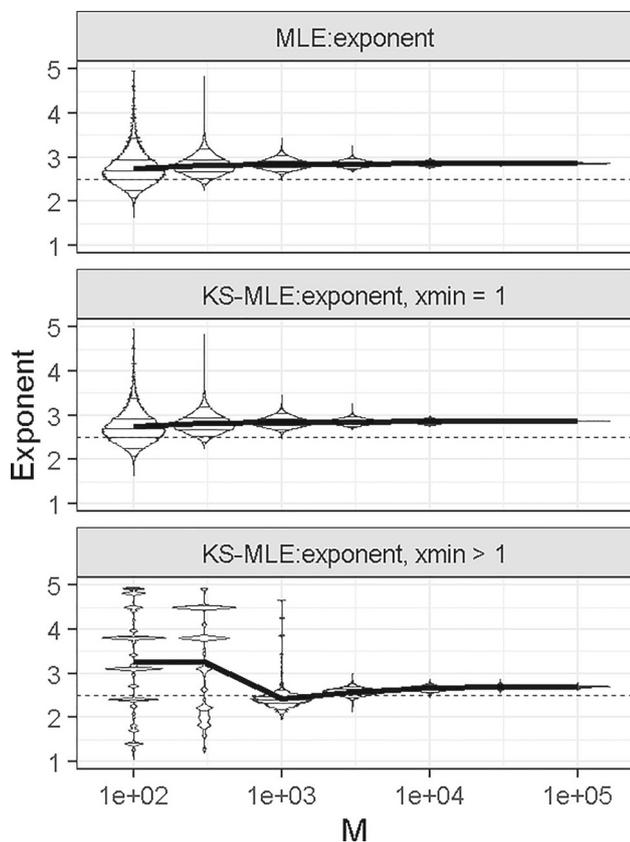


Fig. 3 Coalescence and fragmentation exponent summary statistics for varying M . In all plots, $\bar{F} = 0.20$. Thick black lines connect the means of each violin plot and the thin dashed line is the theoretical value, 2.5. Violin plots are otherwise as in Figure 2

Furthermore, this bias is not eliminated when controlling for the exponential cut-off in the theoretical distribution, nor the finite range of the simulated system (see appendix).

3 Analysis of the base model

Now that we have described our summary statistics and their expected behaviours, we turn our attention to how varying the base model’s underlying rules affects the overall behaviour of the system. We begin by discussing aspects of the basic construction of the model, including different types of kernel, and for the full range of the controlling parameters.

We then discuss in Sect. 4.1 what happens to a particular case of the base model with multiplicative kernel, which has a clear steady state with no gel-shatter cyclicity, when we include various secondary processes. These processes act so as to redistribute small amounts of mass through the system; the effects we observe establish that the system is sensitive to behaviour in the bulk. In Sects. 4.2 and 4.3 we consider the behaviour of the fragmentation function. With these two cases, we describe how the movement of mass from the (right) tail into the bulk influences the behaviour of the system. We summarize these findings in Sect. 4.4 in anticipation of analogous analysis for the cyclic regime.

Size-biased coalescence and shattering fragmentation are clearly not the only possible kernels for the base model, but they are the most intuitive to motivate for most applications. For example, Smoluchowski’s original work used a constant coalescence kernel $K(i, j) = 1$ (without fragmentation) [29], and results for constant, additive, and multiplicative (size-biased) kernels have been known since at least the 1960s for both discrete and continuous mass-distributions. An accessible introduction is [13]. Given the variety of kernels and underlying rules they represent, we first establish some of the sensitivity of the model to the form of the rules, and follow with a brief review of the sensitivity of the model to variation of the parameters, following our previous work [9]. Together, these set the stage for us to robustly interpret further modifications to the system.

First, we consider the significance of the kernels in the system. It is well-known that systems with multiplicative coalescence kernels and, more generally, homo-

geneous kernels exhibit gelation (defined in infinite systems as the divergence of the second moment, [13], p10). Recall that equations (1) track the density n_k and size k of clusters within the system. One could then calculate the mass, $M = \sum_k kn_k$, the total number of units, among which kn_k units have coalesced into the n_k units of size k . For purely coalescent systems (*i.e.* $F(i, k) = 0$) with kernels that exhibit gelation, it is possible for the system to coalesce so rapidly that mass escapes to a cluster of infinite size in finite time, in which case the infinite-sized cluster is referred to as a gel. Typically the remainder of the system is referred to as the sol. In this case, analytically the gel is not tracked by the size-biased sum of the system cluster sizes, resulting in the mass M appearing to decrease in the system even if the equations otherwise appear to conserve mass [13]. In coalescence of random graphs, the gel appears spontaneously as a giant connected component of the system, prior to which the system is made up of isolated graphs. Its formation marks a shift in the size of the largest component; for multiplicative coalescence kernels $K(i, j) \propto ij$ this shift is from order $\log M$ to order M [30]. In our systems (where time is measured by the number of events that have occurred thus far, similar to the number of edges that have appeared thus far in a random graph [30]), ‘gelation’ appears as the creation of a large cluster that rapidly increases in size. This simulated gel (of size k_{\max}) cannot become truly infinite-sized in our simulation, but still separates from the sol, in the sense that the gap between the size of the gel and the next largest cluster is much larger than the next such gap.

It is possible for fragmentation to prevent gelation, but the effectiveness must depend on the relative strength of the fragmentation and coalescence kernels. For example, employing a fragmentation barrier, above which clusters must fragment and restock the smaller cluster sizes, leads to varying steady-state power-law distributions for the sizes of clusters dependent on how fragments and coalescence occur [12]. For certain choices of kernel functions, exact steady-state solutions are possible and well-known, including when the reactions are reversible and satisfy detailed balance [31], such as growth and decay by monomers only (in the presence of sub-linear rate kernels and finite system size) [32]. Alternatively, the steady state can also be reached if a constant source of small clusters is provided to the system to counter-balance the removal by the gel [33]. For coalescence kernels of the form $K(i, j) = \hat{K}(ij)^\alpha$ for $\alpha \geq 0$ and fragmentation of clusters to monomers, a steady-state solution is known to be approximately $k^\alpha n_k \approx \frac{1}{\sqrt{\pi}} n_1 (k+1)^{-3/2}$ and is known exactly for $F(i, k) \propto i$ with $\alpha = 1$ [16, 20, 34, 35]. Many extensions and variants exist, such as non-binary coalescence [21]. Despite the existence of steady-state solutions, however, it is not given that fragmentation actually prevents gelation, although conditions do exist for some forms of fragmentation, such as binary fragmentation [36].

So what happens when systems with shattering fragmentation evolve through time? In Fig. 5, we plot (normalized by the total mass M) the number of clusters against the size of the largest cluster, explored over 20,000 steps in four different simulations, each choosing coalescence or fragmentation to be either multiplicative, as above in equations (2, 3), or constant,

$$K_{\text{const}}(i, j) = \frac{\hat{K}}{N^2}, \quad (5)$$

and

$$F_{\text{const}}(i, k) = \frac{\hat{F}}{N} \delta_{k,1}. \quad (6)$$

(Note that for this constant kernel, in which groups are chosen uniformly at random, normalization is by the total number of groups N rather than the total number of monomers M used to normalize the multiplicative kernel in (2, 3). As there, this reflects a probabilistic interpretation: choose a cluster uniformly at random from the entire population of clusters).

The top two plots use constant coalescence and the bottom two plots multiplicative coalescence, while the left-hand plots use constant fragmentation and the right-hand plots multiplicative fragmentation. We begin with the top-left panel, in which both coalescence and fragmentation are constant, K_{const} and F_{const} . The system aggregates around 2% of its total mass into a single cluster, but the variation is large. Constant kernels act more slowly than, but similarly to, their multiplicative equivalents. The bottom-right panel, in which both coalescence and fragmentation are multiplicative, yields a mostly disaggregated system with some occasional large clusters. The bottom-left panel has multiplicative coalescence K_{mult} but constant fragmentation F_{const} . The largest cluster grows rapidly until the only possible event is fragmentation, at which point the entire system is shattered to clusters of unit size. This behaviour is similar to, but more extreme than, the behaviour we will explore in Sect. 5. On the other hand, when the fragmentation kernel is multiplicative but the coalescence kernel is constant, in the top-right panel, the system never begins to aggregate beyond very small values.

Figure 5 establishes that the form and balance of the kernels are important determinants of the system dynamics. We now restrict to multiplicative kernels, determining how the parameters influence the dynamics of the system, and extending the brief treatment in [9]. We have already noted the existence of stochastic gel-shatter cycles as well as a steady state for coalescence and fragmentation systems. Fixing the scaling of the coalescence and fragmentation kernels, the transition between regimes is determined by the interplay between M , \hat{K} , and \hat{F} . In particular, the dimensionless parameter

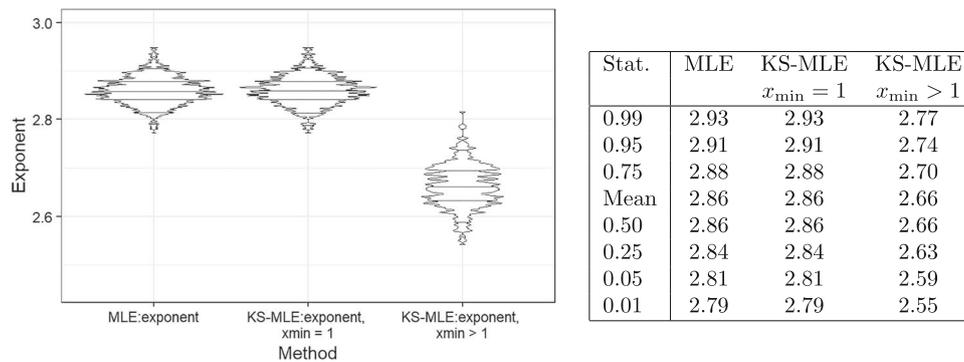


Fig. 4 Coalescence and fragmentation exponent simulated summary statistics. We present the distribution of summary statistics of 1,000 samples from the theoretical coalescence and fragmentation solution with $M = 10^4$ and $\hat{F} = 0.20$. As in Fig. 2, data are presented as violin plots, bandwidth adjustment of 1/8th, for the fitted exponents according to each method (left) with the quantiles (horizontal lines on violin plots) and means (right). Compare with Fig. 2 for which we observe good agreement, indicating deviations from the expected exponent $\alpha = 2.5$ are systematic

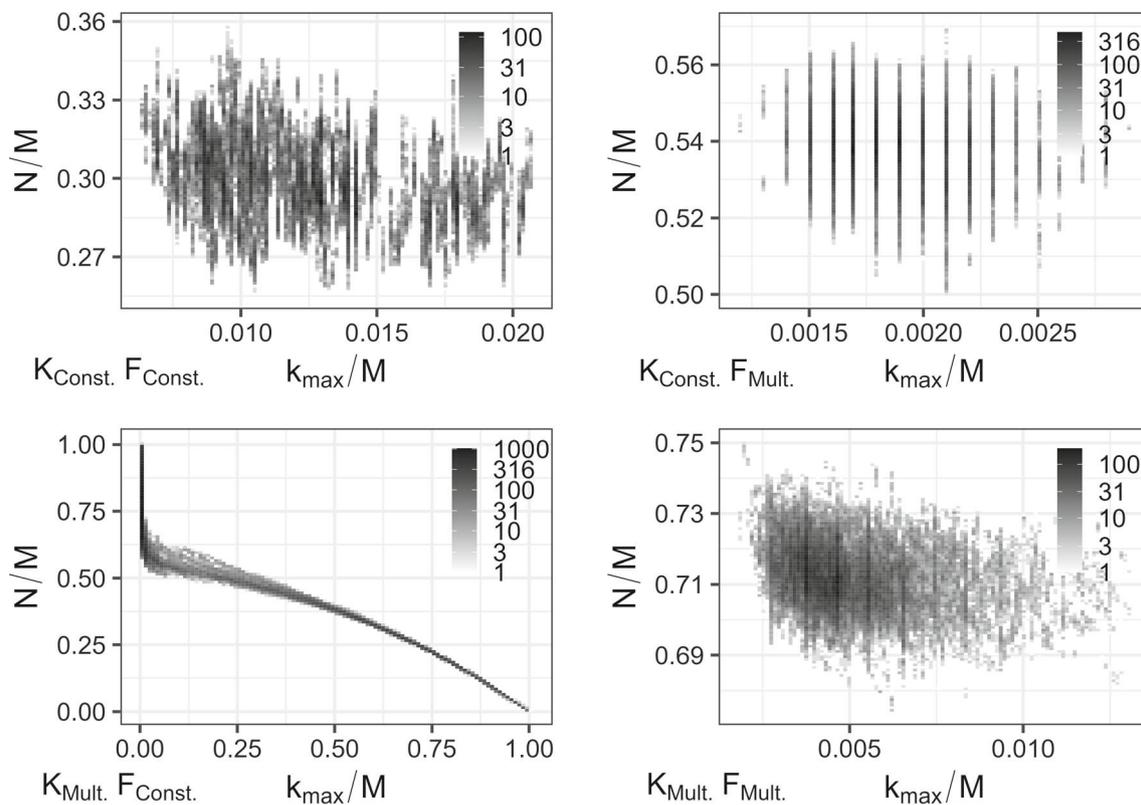


Fig. 5 Space explored by coalescence and fragmentation simulations with varying kernels, $M = 10^4$, $\hat{F} = 0.30$. Top-left: constant coalescence and fragmentation kernels, Eq. 5 and 6. Top-right: we increase the strength of fragmentation by using a multiplicative fragmentation kernel, Eq. 3. Bottom-left: multiplicative coalescence, Eq. 2, constant fragmentation. Bottom-right: multiplicative kernels for both coalescence and fragmentation. When the type of kernel is the same, the system approaches a non-trivial steady state. When it is instead unbalanced the system either remains mostly disaggregated (top-right) or is forced into cyclicity by rapid coalescence followed by fragmentation (bottom-left)

$$r = \frac{\hat{F}M}{\hat{K}} \quad (7)$$

is the ratio of characteristic times of gelation to shattering, and thus captures the balance between the two processes.

When r is small, $r < 0.1$, the system is dominated by rapid gelation and must wait until fragmentation resets the system, similar to when the coalescence kernel is multiplicative and the fragmentation kernel is constant. We expect the cyclicity order parameter \mathcal{K} to be low, $\mathcal{K} < 0.1$, in this regime, as the system rapidly gels and then stalls until a fragmentation event finally occurs. When r is large, $r > 10^3$, the system is instead dominated by fragmentation. When a gel does emerge, its size relative to that of the other clusters and the frequency of fragmentation rapidly removes the gel from the system, forcing the system to dwell primarily in the pre-gel phase. As r increases further, the system begins to resemble the case with K_{const} and F_{mult} . Again, \mathcal{K} is expected to be low in this region, $\mathcal{K} < 0.1$, as the coalescence routinely creates small clusters that are nearly immediately shattered and, otherwise, shattering fragmentation events do nothing to clusters of size 1.

In between the two regimes we find stochastic gel-shatter cycles, the subject of Sect. 5 and previous work [9]. Here, the system visits a broad region of the (k_{max}, N) plot due to the presence of gelation and infrequent but not rare fragmentation. Fragmentation occurs often enough to prevent the system from stalling with all of the mass in a single cluster, but not so often as to prevent the formation of the gel itself. This means that there are many small step increases in the size of the largest cluster alongside a few large step decreases, resulting in a comparatively large \mathcal{K} . These three regimes are contrasted in Fig. 6.

As we have shown in this section, coalescence and fragmentation systems have many built-in modelling assumptions that determine their behaviour. The underlying rules and parameter values naturally have broad impacts, and determine whether and where the system's steady-state behaviour occurs. Of prime importance is the type of kernels used and whether they admit the formation of large clusters or not. If they do, then the balance of the competition between coalescence and fragmentation needs to be correct; too much of either, due to parameters or kernels, and the system stalls.

4 Effects of rule variations on base model with steady state

Having established the influence of kernel type on the system, and as a means to gauge the importance of unmodelled small systematic perturbations in a system of interest, we next consider the effects of process variation by inclusion of a simpler set of *additional* rules in the particular multiplicative-kernel base model case

of Sect. 2, which has $M = 10^4$ and $\hat{F} = 0.2$. This case has an approximate power-law steady state with α confidence intervals KS-MLE [2.568, 2.754] and MLE [2.788, 2.932]. The model is near the edge of cyclicity without fully engaging with stochastic gel-shatter cycles, $r = 2.5 \times 10^3$ and $\mathcal{K} = 0.123$. (See Fig. 1 for the behaviour of this system.)

First in Sect. 4.1 we redistribute small amounts of mass through the system, establishing that it is sensitive to behaviour in the bulk. Then in Sects. 4.2 and 4.3 we consider the behaviour of the fragmentation function, describing how the movement of mass from the (right) tail into the bulk influences the behaviour of the system. We summarize in Sect. 4.4 to prepare ourselves for the cyclic regime.

4.1 Becker–Döring dynamics

The first variation we consider, which we refer to as Becker–Döring dynamics [18, 32, 37], involves only the movement of monomers between clusters in the population. We refer to the addition of monomers to randomly-chosen clusters as ‘accretion’ and the removal of monomers from randomly-chosen clusters as ‘erosion’. These dynamics are motivated by the importance of individuals in social applications such as warfare modelling, although the terminology is that of the physical sciences – accretion might be simply termed ‘recruitment’, while erosion is an individual leaving, through death or disillusionment for example. Our implementation performs fixed numbers n_{ac} of accretion and n_{er} of erosion events each time step (either 0, 1, 3 or 9), and we choose which clusters to accrete or erode at random with probability proportional to their size. Mass remains conserved in our implementation: accretion does not occur if there are no uncoalesced monomers in the system, and erosion does not occur if there are no non-monomers in the system. We note, however, that MLE methods are expected to be more susceptible to perturbations that affect n_1 while KS-MLE should be more robust. We write $\alpha_{(n_{\text{ac}}, n_{\text{er}})}$ and $\mathcal{K}_{(n_{\text{ac}}, n_{\text{er}})}$.

Accretion results in an increase in \mathcal{K} and a decrease in α : $\mathcal{K}_{(9,0)} = 0.831$, mean KS-MLE $\alpha_{(9,0)} = 2.33$, and mean MLE $\alpha_{(9,0)} = 1.72$. More comparable to our earlier statistics, we have $\mathcal{K}_{(3,0)} = 0.564$, KS-MLE $\alpha_{(3,0)} \in [2.248, 2.682]$, and MLE $\alpha_{(3,0)} \in [1.906, 2.136]$ in 95% of samples. These reflect the increase in number of larger clusters and the tendency to increase the largest cluster's size even when the coalescence events do not target this cluster. The increase in \mathcal{K} is somewhat deceptive, as can be seen in Fig. 7 (top-left), especially in comparison to Fig. 1 and the system studied in Sect. 5, Fig. 10. The system is certainly more cycle-like, given the shape of the state space visited has elongated, but it has not fully escaped steady state. Accretion shows that the system's supposed power-law can be heavily weighted towards larger clusters by increasing the recruitment amongst monomers only.

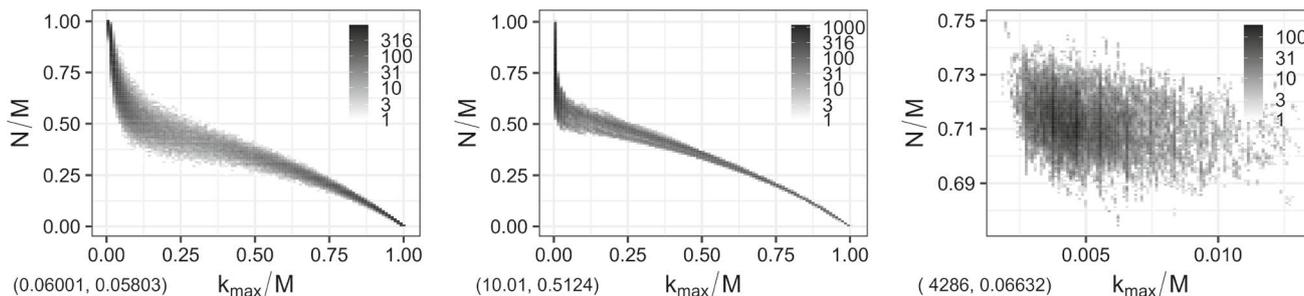


Fig. 6 Shapes of spaces explored in various regimes, kernels K_{mult} and F_{mult} . We vary simulation parameters M and \hat{F} to showcase the different regimes, obtaining different r and \mathcal{K} , shown in parentheses at bottom left of each panel. In the first plot with $M = 300$ and $\hat{F} = 0.0002$, we obtain forced fragmentation cycles, in which the system quickly coalesces into a single cluster (bottom right of plot) and then must wait for a rare fragmentation event to reset it. In the second plot with $M = 10000$ and $\hat{F} = 0.001$, the system continuously grows and fragmentation is stochastic and common, preventing stalling and creating unforced stochastic gel-shatter cycles. In the third plot, $M = 10000$ and $\hat{F} = 0.3$, fragmentation dominates, preventing the emergence of a gel. See [9] for further details

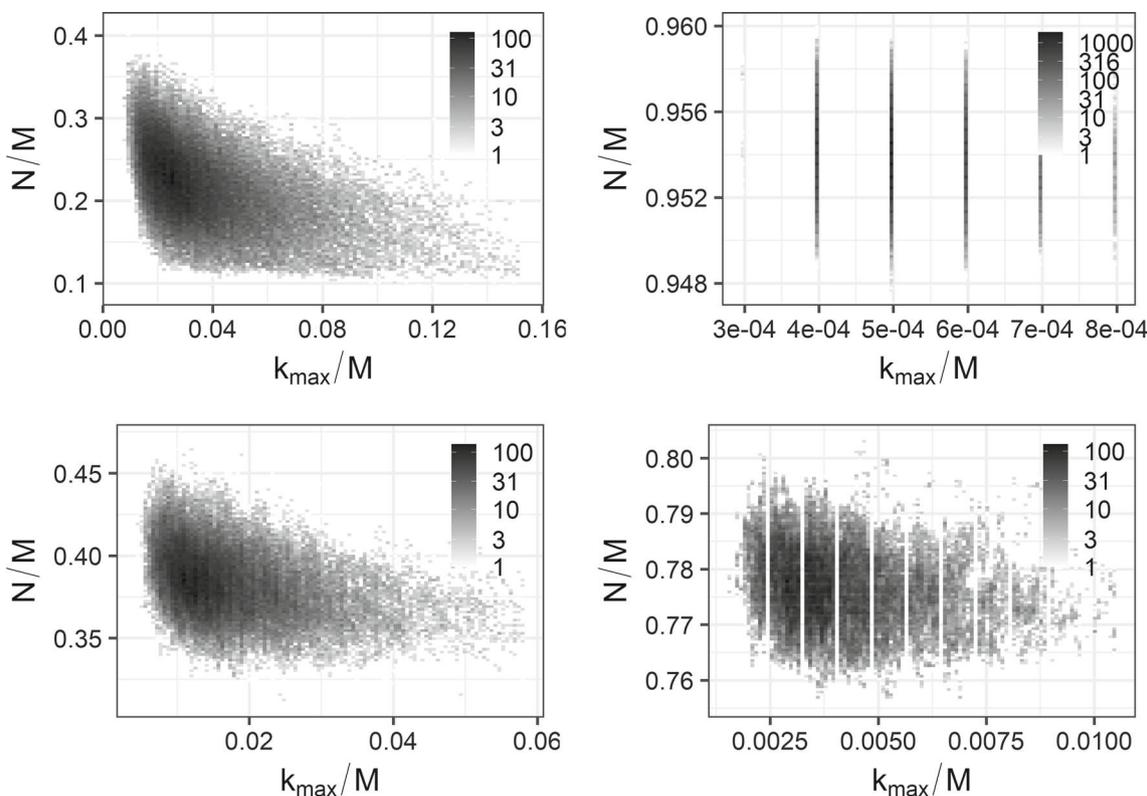


Fig. 7 Clockwise from top-left: accretion, erosion amongst any clusters, erosion amongst unique clusters, and combined accretion and erosion. In each case, we set $n_{ac} = 9$ and $n_{er} = 9$ to best show the effects

Erosion generally shows the opposite behaviour to accretion, but its effectiveness relies on an implementation detail: do we remove from distinct clusters within a time step, or can the same cluster have multiple removals? To do the former greatly restricts the effects of erosion, as can be seen in Fig. 7 (top-right and bottom-right) where $n_{er} = 9$, and still results in a steady-state-like system. The summary statistics are $\mathcal{K}_{(0,3)} = 0.007$, KS-MLE $\alpha_{(0,3)} \in [2.849, 3.047]$, and MLE $\alpha_{(0,3)} \in [3.122, 3.229]$ in 95% of samples. To do the

latter prevents the system from growing beyond $k_{max} = 8$, and a power-law is a poor description of this case with $n_{er} = 9$. When $n_{er} = 3$, the summary statistics are $\mathcal{K}_{(0,3)} = -0.011$, KS-MLE $\alpha_{(0,3)} \in [2.118, 4.848]$, and MLE $\alpha_{(0,3)} \in [3.528, 3.635]$ in 95% of samples.

Overall, the system is less robust to erosion, the frequent loss of monomers from clusters, than to accretion, frequent gain of monomers by clusters. The exact implementation can have a drastic effect, with more frequent

erosion reducing cyclicity and making the system more steady-state-like.

For completion, combining accretion and erosion with $n_{ac} = n_{er}$ results in a hybrid system that is similar to but more aggregated than the original system, Fig. 7 (bottom-left). There is a modest increase in \mathcal{K} : $\mathcal{K}_{(1,1)} = 0.148$, $\mathcal{K}_{(3,3)} = 0.179$, and $\mathcal{K}_{(9,9)} = 0.215$ when using erosion without the uniqueness restriction. This reflects a slightly more cyclic system that is still not far from steady state. The system also suffers from a depletion of monomers, so while KS-MLE $\alpha_{(3,3)} \in [2.468, 2.806]$, MLE $\alpha_{(3,3)} \in [2.178, 2.315]$.

Finally we note that a novel phenomenon, of super-clustering states, has recently been observed in size-biased generalizations of the above Becker-Döring model. Such states are characterized by intensive fluctuations which violate the standard van Kampen expansion [38, 39].

4.2 Weakening fragmentation

We have learned that accretion and erosion can be quite influential, but can also balance each other (resulting in a generally more aggregated system without greatly affecting the gelation-related properties) and generally make coalescence and fragmentation work more efficiently. Against this, the models' assumptions on coalescence and fragmentation already take forms that are probably more efficient than in real systems. (For example, collisions might join two clusters imperfectly, producing some fragments; or shattering may fail to produce only monomers.) Overall this prompts the question, how efficient must the system be to be well-approximated by shattering fragmentation? (Coalescence has been treated elsewhere in the literature, *e.g.* [11].) For example, in an insurgency war situation complete shattering could be interpreted as a perfect response of a cell to being compromised – idealized but implausible.

To attempt to answer this question, we consider a suite of variations of fragmentation. We begin with halving a cluster (*i.e.* a size- n cluster fragments into one of size $\lfloor \frac{n}{2} \rfloor$ and one of size $\lceil \frac{n}{2} \rceil$ where $\lfloor \cdot \rfloor$ and $\lceil \cdot \rceil$ are the floor and ceiling functions). We then consider the case where one of these two clusters is shattered, before altering the proportion of the original cluster shattered from one-half to one-tenth or nine-tenths. We will see that even a small amount of shattering can significantly reset a system while keeping it near steady state. Next we create the fragments using a single partition drawn uniformly at random from the set of possible partitions of the original cluster. This motivates us finally to consider a specific intermediate possibility, repeated fragmentation of a cluster or *stick-breaking*, in Sect. 4.3.

Halving a cluster instead of shattering it is perhaps the furthest removed from shattering fragmentation: not only is it reversible in a single (binary coalescence) step, but it leaves both newly formed clusters as aggregated entities. It is perhaps no surprise then that coalescence and such a weak form of fragmentation leave the

system in a very small number of clearly discrete states, as seen in Fig. 8 (top-left). We find $\mathcal{K} = -0.005$; the system spends slightly more time halving the largest group than it does growing it. To apply power-law distribution estimators would be clearly incorrect here, even as proxies for other processes. If we attempt to do so, approximately 79% of KS-MLE fits ($x_{\min} > 1$) report $\alpha = 0$.

Instead, consider the case where we shatter some portion of the original cluster to be fragmented. This still leaves the remainder as an aggregated cluster, but requires more – potentially many more – coalescent steps to reverse the damage caused by the fragmentation. In Fig. 8 we consider shattering 0.5 of the cluster (top-right), 0.1 of the cluster (weakening fragmentation, bottom-left), and 0.9 of the cluster (strengthening fragmentation, bottom-right). Each result looks consistently steady-state-like (compare Fig. 1, top-left), but we are also seeing a smooth but modest deformation in the shapes of the plots (note that the scales on the horizontal axes differ), with the largest clusters tending to be smaller as the extent of shattering increases. This is reflected in the summary statistics. Shattering 10% of a cluster results in 95% of KS-MLE α estimates in $[2.491, 2.721]$ and MLE α estimates in $[2.759, 2.953]$. Shattering 50% yields estimates of $[2.431, 2.659]$ and $[2.727, 2.875]$ respectively, while shattering 90% yields $[2.548, 2.744]$ and $[2.783, 2.918]$. Compared to the original coalescence and fragmentation estimates of $[2.568, 2.754]$ and $[2.788, 2.932]$, the differences are quite small. While the occupied region of the $(\frac{k_{\max}}{M}, \frac{N}{M})$ plane is somewhat different, it seems that even relatively small amounts of shattering can fairly well approximate full shattering for some purposes. That said, increasing shattering has a pronounced effect on the tendency of the system to produce stochastic cyclic or cycle-like behaviour. Shattering 10%, 50% and 90% of the cluster yield $\mathcal{K} = 0.431, 0.186,$ and 0.126 respectively. The summary statistic \mathcal{K} of cyclicity is thus more affected by the proportion shattered than is the α of the steady state, with highly inefficient shattering affecting the power law very little but creating much greater stochastic cyclicity hidden behind it.

So far, coalescence and fragmentation appear to be satisfactorily robust to perturbations to the exact form of the fragmentation process, as long as some shattering is present. But once again this may be an idealization motivated by the physical sciences – in the social sciences it could be that fragmenting is more typically into a range of group sizes, and complete return to individual autonomy is implausible. So, does the system remain robust when we move away from shattering? And how do the sizes of the clusters formed complicate the situation?

There are many ways to fragment a cluster, a problem which is equivalent to partitioning a natural number. We begin first with ‘probabilistic divide and conquer’, which selects a partition uniformly at random from the set of partitions [40]. Our implementation results on average in 25 monomers (95%: $[0, 88]$) and 94.6 clusters (95%: $[53, 165]$) when partitioning 1,000 in 1,000 trials,

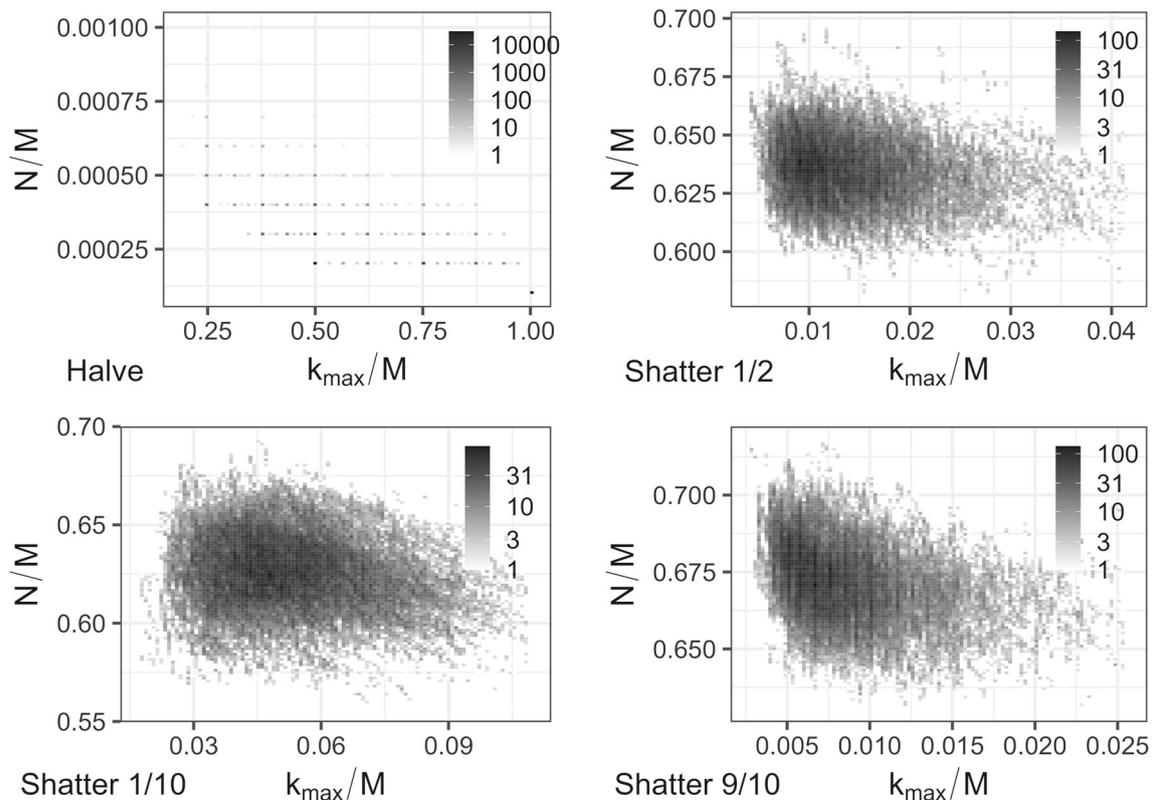


Fig. 8 Clockwise from top-left: halving fragmentation, shattering half fragmentation, shattering nine-tenths fragmentation, and shattering one-tenth fragmentation

so only slightly fewer clusters, but the clusters are larger in the median than for shattering one-tenth. (Note that shattering one-tenth would result in 101 clusters, with 100 monomers, shattering one-half 501 or 500, and shattering nine-tenths 901 or 900.) Despite the similar number of clusters, the exact size of clusters is important. Partitioning fragmentation results in 95% of KS-MLE α estimates in $[2.22, 2.62]$ and 95% of MLE α estimates of $[2.09, 2.19]$. Furthermore, the space itself is far more aggregated on average than that of one-tenth shattering fragmentation. While one-tenth shattering fragmentation had around 6254 clusters on average, partitioning fragmentation has around 4032, although this is counterbalanced by a higher average x_{\max} for one-tenth shattering fragmentation (532 vs 204).

So we now know that the exact form of the fragmentation does matter, and that even shattering just one-tenth of a cluster makes the system fairly close to that in which the entire cluster is shattered. On the other hand, this seems to be serving as a proxy for the ease of re-assembly of a fragmented cluster. Due to the size-biased nature of the kernels, obtaining monomers makes it harder to re-assemble than if one obtains clusters of varying sizes. Unfortunately, partitioning uniformly at random does not give us a clear control parameter to explore how the results are changing. In the next section, we explore a method of partitioning that does have such a parameter, allowing us to explore its effect on our summary statistics.

4.3 Stick-breaking and Chinese restaurants

One simple form of fragmentation, analogous to stick-breaking processes [41], is to repeatedly fragment a portion of a cluster. We take a stick (respectively, cluster), snap it in two at a random point (resp. divide it into two parts), discard one, and then repeat the process on the remainder. The question then becomes how to pick our random point. We begin here with a uniform distribution, before considering the natural extension to the beta-binomial distribution. We then consider a more general parametrized form of fragmentation, the Chinese Restaurant Process (CRP).

We can also use stick-breaking to create some coalescence, although a little more specification is needed. When two clusters meet, our implementation transfers to the larger cluster the first fragment removed from the smaller cluster, which must be of size at least 1, after which the remainder of the smaller cluster experiences stick-breaking fragmentation as described above. See appendix.

Figure 9 contains the results of stick-breaking coalescence (with base model fragmentation; top-left) and stick-breaking fragmentation (with base-model coalescence; top-right). Initial inspection of the former might lead one to believe that stick-breaking coalescence achieved no change from the base case other than a slightly damped cyclicity, $\mathcal{K} = 0.094$ from 0.123 and a slightly lowered maximum group size, with 95% of

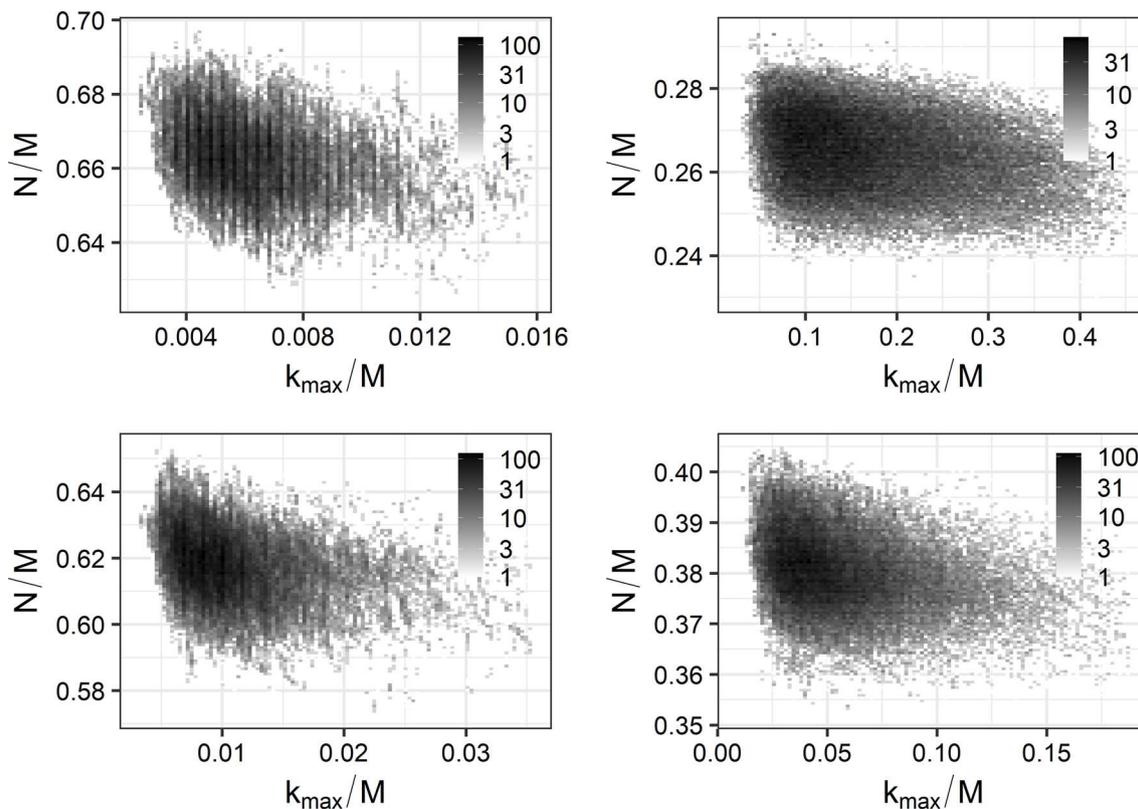


Fig. 9 Clockwise from top-left: stick-breaking coalescence, stick-breaking fragmentation, CRP fragmentation with $\theta = 1.2$, and CRP fragmentation with $\theta = 1.8$

KS-MLE and MLE estimates of $\alpha \in [2.534, 2.768]$ and $[2.725, 2.829]$ respectively. We will return to this case in a moment. On the other hand, stick-breaking fragmentation causes a larger perturbation to the system than does partitioning fragmentation, with $\mathcal{K} = 0.412$ and 95% of KS-MLE and MLE estimates of $\alpha \in [1.890, 2.166]$ and $[2.241, 2.402]$ respectively. Notice that whereas partitioning fragmentation reduced the MLE more, stick-breaking fragmentation reduces the KS-MLE more. This suggests that the tail induced by stick-breaking is far heavier than that produced by partitioning. In contrast, partitioning has a larger effect on the number of monomers in the system.

Finally, the combination of stick-breaking coalescence and fragmentation helps reveal some of the unseen effects of stick-breaking coalescence. While not entirely capable of reining in the effects of stick-breaking fragmentation, stick-breaking coalescence has a large impact, bringing \mathcal{K} to 0.236 and 95% of KS-MLE and MLE estimates to $\alpha \in [2.208, 2.393]$ and $[2.298, 2.369]$. Stick-breaking coalescence permits a modest decrease in the speed of gelation while making the system less aggregated in general. This in turn makes it harder to reverse the effects of a single stick-breaking fragmentation because there is a smaller reservoir of medium-large clusters (in exchange for a larger reservoir of small-medium clusters). It also appears that this stick-breaking coalescence does not have a large effect on the α estimates by itself because the

exchange of small-medium clusters for medium-large clusters does not greatly influence either the tail (KS-MLE) or the number of monomers (MLE).

Moving to a beta-binomial distribution, thus allowing the individual breaks to be more biased towards larger or smaller clusters, does not change the results greatly. Moving the distribution's parameters from $(1, 1)$ (which is the specialization to the uniform distribution) to $(2, 4)$, $(3, 3)$, and $(4, 2)$ results in 95% KS-MLE α estimates of $[2.044, 2.221]$, $[1.941, 2.151]$ and $[1.853, 2.147]$ respectively.

Another interpretation of discrete stick-breaking gives us smooth control of fragmentation: the Chinese Restaurant Process (CRP). The CRP is named for its metaphorical construction of a partition. For a partition of size n , we consider a queue of n people who are permitted to seat themselves one at a time at circular tables. After the first individual is seated, the second then has a choice to sit with the first, or to initiate a new table, and so on for each customer. Further, it is presumed that the probability that a new customer picks an occupied table is proportional to the number of people sitting there. This mechanistic process has natural analogues to the manner in which a sub-group in insurgency warfare might fragment under external pressure. The CRP has two parameters, which when positive can be interpreted as controlling how intrinsically attractive it is to start a new table (the 'strength') and slightly penalizing tables already occupied (the 'discount'). Of these

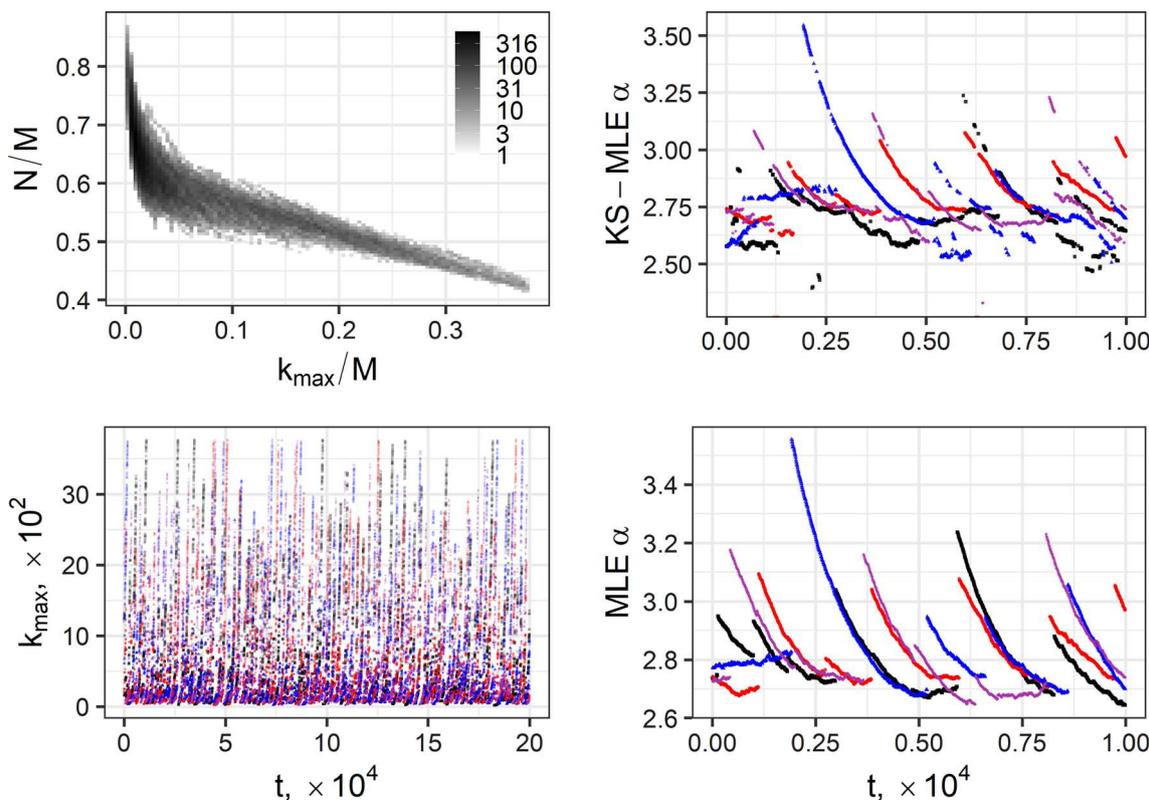


Fig. 10 Coalescence and fragmentation summary statistics, $M = 10^4$, $\text{Pr}(\text{Frag.}) = 0.01$. Plot types are the same as in Fig. 1. The time series shown for estimates of the exponent are shorter than those for k_{\max} or those in Fig. 1 to better show details

two, the latter, say $\phi \in [0, 1]$, is more important for our purposes, as it controls the power-law distribution over the number of customers seated at each table with exponent $\theta = 1 + \phi$ [41, 42]. (In principle, power laws above $\theta = 2$ could be accessed via *e.g.* Price’s network model [43], but such power laws did not appear to substantially alter our results.)

Replacing our stick-breaking fragmentation with CRP fragmentation for various θ proves particularly tractable for navigating various summary statistics. Cyclicity \mathcal{K} now appears to vary smoothly: compare stick-breaking fragmentation’s $\mathcal{K} = 0.412$ with $\theta = 1.2$ yielding $\mathcal{K} = 0.436$, $\theta = 1.5$ $\mathcal{K} = 0.246$, and $\theta = 1.8$ $\mathcal{K} = 0.156$. Similarly α estimates by both KS-MLE and MLE methods proceed smoothly: our 95% intervals for KS-MLE α go from $[1.890, 2.166]$, through $[2.075, 2.257]$ and $[2.299, 2.426]$, to $[2.459, 2.620]$, while MLE α proceeds as $[2.241, 2.402]$, $[2.274, 2.376]$, $[2.434, 2.511]$, and $[2.640, 2.734]$. This is reflected in the bottom panels of Fig. 9, which show results for $\theta = 1.2$ and $\theta = 1.8$. Just as in Fig. 8, the system appears to be returning to the original steady state as we go from stick-breaking fragmentation (top-right; largest k_{\max}/M and smallest N/M) through CRP with $\theta = 1.2$ (bottom-right; intermediate values) to CRP with $\theta = 1.8$ (bottom-left; smallest k_{\max}/M and largest N/M).

It is notable that one power-law distribution producing process has such a strong and well controlled

effect when used as a part of a wider power-law distribution producing process. Presumably, the reason why this works so well is, building on the above, a power-law distribution with exponent θ in $[1, 2]$ is not so clustered as to have no effect (mostly shattered, $\theta > 2$) while providing a consistent control on shape. This is in contrast to beta-binomial which clumps too much, while a significant amount of shattering takes too long to reverse (since monomers are slow and difficult to coalesce back together). These results are supported by previous work by Brilliantov et al. [34], which established that if the power-law distribution of fragments is steep enough, *i.e.* $\theta > 2$, the system is close to complete shattering of clusters. Notably, these results were for fragmentation caused by (binary) collisions, suggesting that our robustness results might hold for a broad class of CF models.

4.4 Summary

Is it safe to model a CF system as being in steady state? In this section, we have considered how important unmodelled systematic perturbations can be to the steady state of a CF system that is believed to be in or close to its steady state.

Often, the CF system proved quite robust. Small amounts of accretion and erosion do not have disproportionate effects and dominate the system. Greater

accretion and erosion can, however, shift the system away from its original power-law steady state. For large amounts of accretion and erosion to cancel out each other's effects would require fine-tuning. Further, shattering does not need to be total to replenish the resupply of monomers: a small amount of shattering effectively replicates the effects of total shattering.

At the extreme of simplicity, replacing shattering with halving altered the system beyond recognition, placing it firmly in a regime of simple forced cycles from which it cannot escape. More subtly, there are distributions of fragments which leave enough medium-sized clusters that the system is delayed from reassembling without forcing it to reset (nearly) completely or not at all. While this can be achieved with random partitioning or varieties of stick-breaking, this regime is most accessible by a power-law distribution generating process such as the Chinese Restaurant Process, a mechanistic process originally motivated by an analogy to people's behaviour. In such circumstances the cyclicity parameter \mathcal{K} can easily reach values significantly greater than zero, indicating some form of time-asymmetric cyclicity and, minimally, that the state is rather unsteady, so that a simple assumption of a steady state described by a power law is certainly not telling the full story.

5 Effects of rule variations on base model with stochastic gel-shatter cycles

As we saw in the last section, it is not guaranteed that a steady state will emerge from a coalescence and fragmentation process. A prominent alternative is the formation of *stochastic gel-shatter cycles*, in which the system is dominated by gelling coalescence and then stochastically resets due to shattering fragmentation. We use a model gel-shatter system to show that this regime is more susceptible to perturbations of the sort described in Sect. 4. For the model system seen in Fig. 10, we set $M = 10^4$ and $\hat{F} = 0.01 = 1 - \hat{K}$ for $r \approx 101$, yielding $\mathcal{K} = 0.490$. This system starts in a very disaggregated state, grows, and then reaches a critical gelation point beyond which it rapidly aggregates before stochastically shattering and repeating. As it does so, the fitted exponent α decreases smoothly before jumping upwards. Its exponent summary statistics are the confidence intervals for KS-MLE α and MLE α from 95% of samples. For this case, they are respectively [2.474, 2.829] and [2.656, 3.163]. We proceed in parallel to Sect. 4: we first show that additional processes can significantly alter the α statistics before again observing that power-law distributed fragmentation greatly affects the stochastic gel-shatter cycles. Summary statistics will be presented in Table 2 at the end of this section.

We begin by adding Becker-Döring mechanics to our base model. The effects of extreme accretion, $n_{ac} = 9$, can be seen in Fig. 11 (top-left). Accretion exagger-

ates the stochastic gel-shatter cycles, greatly increasing \mathcal{K} with $\mathcal{K}_{n_{ac}=1} = 0.623$, $\mathcal{K}_{n_{ac}=3} = 0.721$, and $\mathcal{K}_{n_{ac}=9} = 0.797$ and widening the sampled KS-MLE $\alpha_{n_{ac}=3}$ to [2.261, 3.043]. Naturally, accretion decreases n_1 , more substantially influencing the MLE α , which predicts a flatter distribution than KS-MLE α does. The gel-shatter system is fairly robust to accretion, in the sense that α does not change dramatically and behaviour is similar to, albeit an exaggeration of, the typical behaviour for this regime. Identifying that a system was undergoing accretion as well as coalescence and shattering would be a difficult task as a result, although there are substantial differences in the behaviour of N and k_{max} that might be relevant for specific applications.

Erosion has similar effects on gel-shatter and steady-state systems. When erosion occurs amongst any clusters, it rapidly forces the system to a steady state of only small clusters for $n_{er} = 9$, while requiring erosion amongst separate unique clusters forces the system nearer to but not fully into a steady state, as shown in Fig. 11 (top-right and bottom-right). This is reflected in the differences in summary statistics in Table 2: $\mathcal{K} = -0.079$ and KS-MLE $\alpha \in [2.904, 3.191]$ versus $\mathcal{K} = 0.178$ and KS-MLE $\alpha \in [2.615, 2.991]$ respectively. When erosion is amongst any clusters, this predicts that it is more likely for the system's largest cluster to shrink on any given time-step. In effect, \mathcal{K} suggests that frequent erosion can act as a suitable alternative to fragmentation. In either case, we must conclude that erosion can perturb the expected power-law exponent and even cause a gel-shatter system to behave as a (modified) steady-state system. We conclude our discussion of accretion and erosion with the note that, as before, combining accretion and erosion mostly moderates the two constituents, resulting in somewhat higher \mathcal{K} and depleting monomers but generally similar behaviour as seen in Fig. 11 (bottom-left).

Clearly, simple Becker-Döring dynamics can significantly affect gel-shatter dynamics. It is then not surprising that using power-law variations should also significantly alter the dynamics observed. We focus on power-law distributed CRP fragmentation. At the extreme is stick-breaking fragmentation, seen in Fig. 12 (top-left). In this case, the system remains almost completely aggregated despite experiencing fragmentation. Indeed, $\mathcal{K} = -0.026$ indicates that fragmentation is the dominant process in the system, reflected in the extremely low fitted exponents. It is hard to argue that the system is power-law distributed, however, rendering the exponent summary statistics superfluous. Stick-breaking fragmentation effectively results in a system that is almost always aggregated into a single cluster, which rarely fragments slightly before quickly re-aggregating the fragments. In this sense, stick-breaking fragmentation removes any gel-shatter cyclicity.

Low fragmentation power-law exponents θ reproduce results akin to that of stick-breaking fragmentation. As θ increases, the system becomes more disaggregated after fragmentation and requires longer periods of time to re-aggregate fully. Large clusters are still common

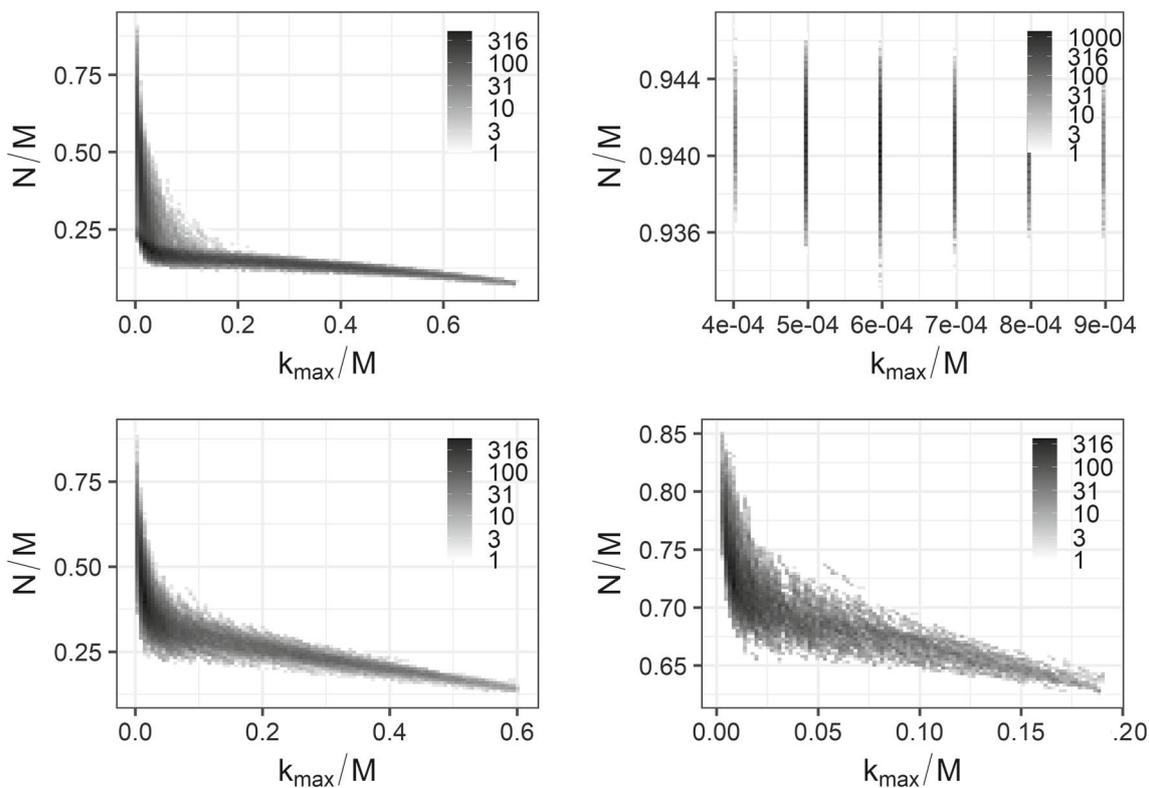


Fig. 11 Clockwise from top-left: accretion, erosion amongst any clusters, erosion amongst unique clusters, and combined accretion and erosion. In each case, we set $n_{ac} = 9$ and $n_{er} = 9$ to best show the effects

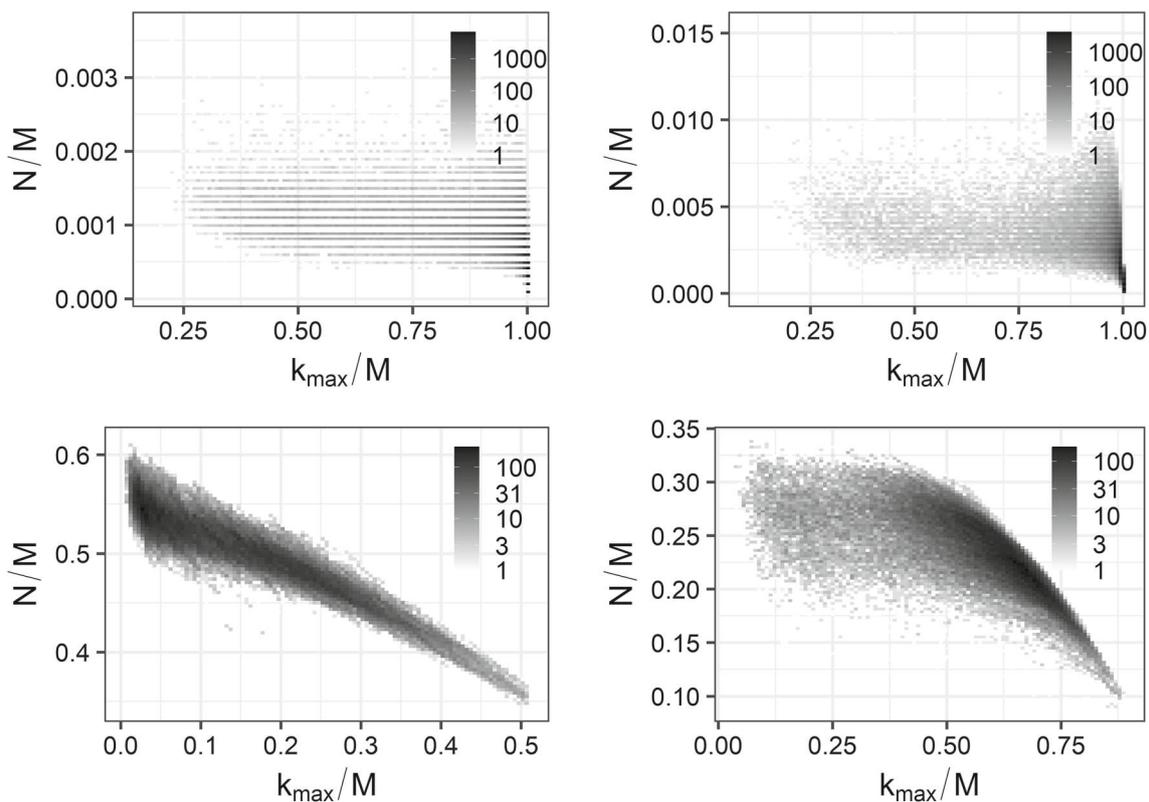


Fig. 12 Clockwise from top-left: stick-breaking fragmentation, CRP fragmentation with $\theta = 1.2$, CRP fragmentation with $\theta = 1.5$, and CRP fragmentation with $\theta = 1.8$

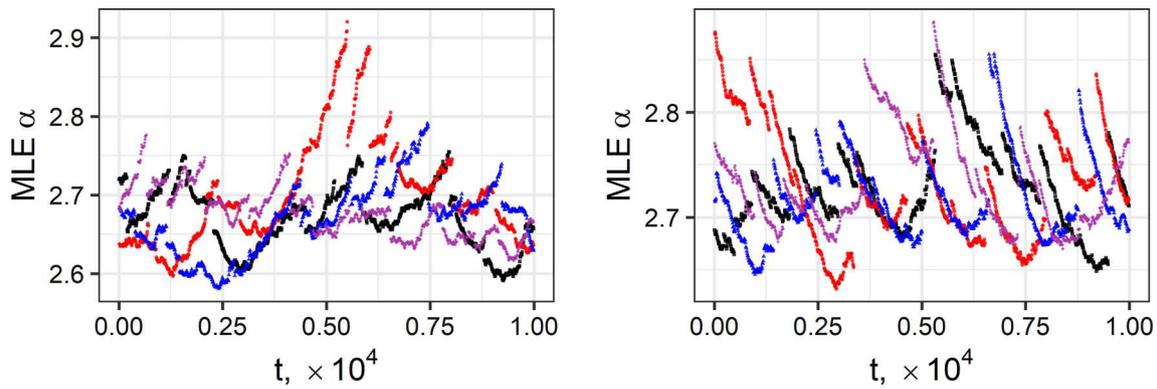


Fig. 13 Coalescence and fragmentation MLE exponent summary statistics combining stochastic gel-shatter cycles with power-law distributed fragmentation, $M = 10^4$, $\text{Pr}(\text{Frag.}) = 0.01$, and $\theta = 1.80$, left, and $\theta = 1.90$, right. Four simulations were conducted (black, blue, red, and purple). Compare with Fig. 10

Table 2 Summary of robustness results in stochastic gel-shatter cycle coalescence and fragmentation. In all cases, $M = 10^4$, $\text{Pr}(\text{Frag.}) = 0.01$, and four simulations were used. The ranges given for KS-MLE and MLE estimates of power-law α contain 95% of empirical results. Perturbations used are explained in the text. $\mathcal{K} > 0.2$ is characteristic of gel-shatter cycling. We exclude large (> 5) and small (< 1) estimates of the KS-MLE α from our intervals

| Perturbation | Cyclicity \mathcal{K} | KS-MLE ($x_{\min} > 1$) | MLE |
|-----------------------------------|-------------------------|---------------------------|----------------|
| Stochastic gel-shatter cycles | 0.490 | [2.474, 2.829] | [2.656, 3.163] |
| Accretion (3, Unique) | 0.729 | [2.248, 2.957] | [1.559, 2.800] |
| Erosion (3, Unique) | 0.178 | [2.615, 2.991] | [2.970, 3.316] |
| Accretion and Erosion (3, Unique) | 0.590 | [2.370, 2.884] | [1.960, 2.710] |
| Accretion (3) | 0.721 | [2.261, 3.043] | [1.603, 2.849] |
| Erosion (3) | -0.079 | [2.904, 3.191] | [3.276, 3.435] |
| Accretion and Erosion (3) | 0.524 | [2.420, 2.883] | [2.062, 2.673] |
| Fragmentation Stick-breaking | -0.026 | [0.686, 1.563] | [0.902, 1.440] |
| CRP $\theta = 1.20$ | 0.151 | [0.979, 1.838] | [1.027, 1.956] |
| CRP $\theta = 1.50$ | 0.910 | [1.875, 2.519] | [2.357, 2.964] |
| CRP $\theta = 1.80$ | 0.811 | [2.362, 2.609] | [2.588, 2.799] |

after fragmentation, aiding the reaggregation of the system by decreasing the number of steps that do not contribute directly to k_{\max} and thus increasing \mathcal{K} beyond that of the standard system: $\mathcal{K}_{\theta=1.2} = 0.151$, $\mathcal{K}_{\theta=1.5} = 0.910$, and $\mathcal{K}_{\theta=1.8} = 0.811$. The system also experiences a smooth change in the fitted power-law exponents as the system transitions back to the original gel-shatter system: stick-breaking has KS-MLE $\alpha \in [0.686, 1.563]$, $\theta = 1.2$ has KS-MLE $\alpha \in [0.979, 1.838]$, $\theta = 1.5$ has KS-MLE $\alpha \in [1.875, 2.519]$, and $\theta = 1.8$ has KS-MLE $\alpha \in [2.362, 2.609]$. Discussions of summary statistics and state spaces do not show the full range of effects of power-law fragmentation, however. Comparing Fig. 13 with the distinctive cycles observed in the base gel-shatter model of Fig. 10, one still sees some cyclicity, but over a smaller range of α , less distinctive, and with greater stochasticity both overall and in the rougher variation in α observed within each cycle. There is also a subtle transition between falling α , Figs. 10 and 13 right, and rising α , Fig. 13 left.

6 Conclusion

Descriptions of coalescence and fragmentation (CF) are ubiquitous across many scientific and social-scientific applications, yet there is limited understanding of the connections between the microscopic processes, theorized or actual, and the observed fragment distributions. There are various specific models, often analytically tractable at the level of deterministic treatment of means, that result in (approximate) power-law distributions of group sizes. Empirical distributions, too, are often observed to be approximately power-law, so that power laws [24] have become the standard, go-to method for describing the outcomes of CF processes [e.g. 12].

But this is not enough on its own to justify a claim that the models are correctly capturing the microscopic processes, or that the power-law-described steady state is the best approximation to the macroscopic statistics implied by these processes. The danger is of a ‘street-light effect’, in which the model is used because it is

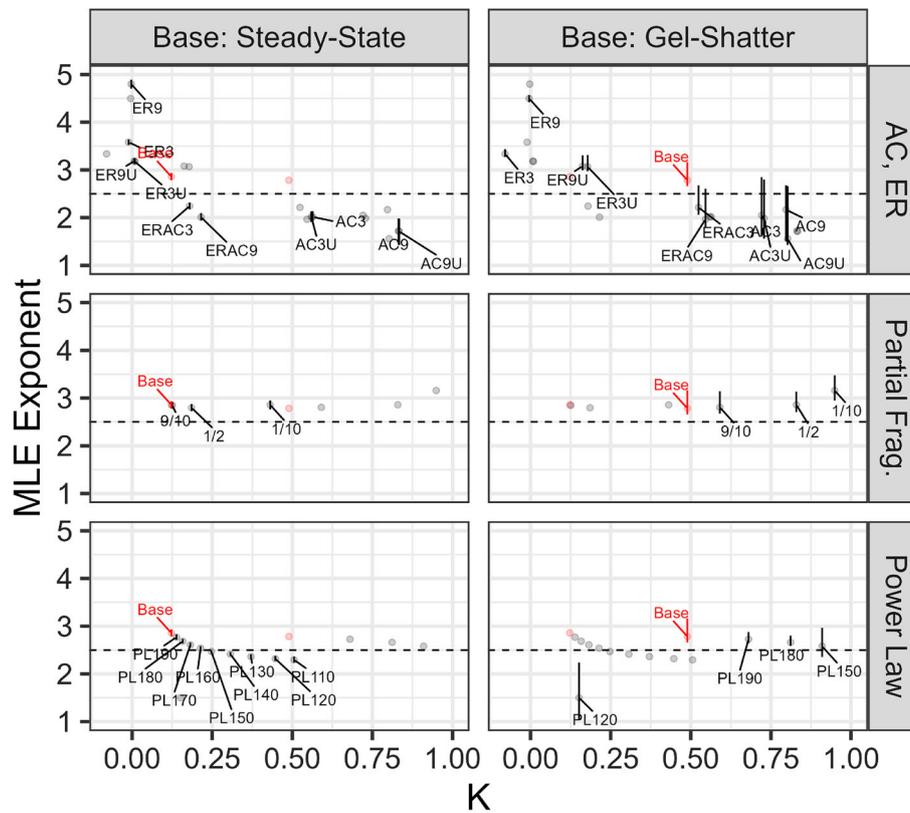


Fig. 14 Visual summary of robustness results for both steady-state and stochastic gel-shatter cycles. Select results from Tables 1 (left column) and 2 (right column) are displayed as MLE fitted exponent against the cyclicality \mathcal{K} . Fitted exponents are displayed as vertical 95% ranges with points placed at the median estimates. The base case is highlighted in red, a horizontal dashed line at the expected 2.5 exponent is plotted (but see Figs. 2 and 4), and the points are repeated between the two base states (*i.e.* within a row) to highlight shifts. Results for KS-MLE are similar, but show less curvature for accretion and erosion. The shorthand labels are: ACn_{ac} , accretion of n_{ac} ; ERn_{er} , erosion of n_{er} ; U, with the uniqueness constraint; 9/40, 1/2, 1/10, shattering of proportion of cluster; $PL\theta$, fragmentation according to a CRP with parameter θ (units of hundredths here)

available, tractable and sufficient to match the empirical statistics, rather than because it is correctly modelling the underlying process.

The purpose of this article was to take a standard model with an approximately power-law steady-state distribution and fully explore the finite-population stochastic simulations of the model with variations of its microscopic rules, observing the effect on the steady state and summary statistics. We examined the maximum likelihood estimators for the power-law exponent α for the whole distribution and (by minimizing the Kolmogorov–Smirnov distance between the actual and fitted cumulative distributions) for the most power-law-like part of the distribution, typically the tail.

Such a task is complicated by the possibility of cyclicity: is a steady state being imputed to a system which, on timescales important in the applied context, is not steady but rather cyclic? – is there no detailed balance? As we saw in a precursor article [9], the phenomenon of gelation (aggregation into a single large cluster), if accompanied by stochastic shattering of clusters, can create stochastic gel-shatter cycles. Not all stochastic cyclicity is necessarily of this form, but any cyclicity will

still be a significant departure from a steady state. The best summary statistic to capture this was a measure \mathcal{K} of time-asymmetric cyclicality, the asymmetry between coalescing and fragmenting time steps.

A full analysis of how any variation in the microscopic details gives rise to power law steady states and/or cyclicity would be a major research programme in itself. Our numerical results, presenting unified data from Tables 1 and 2 in Fig. 14, enable us to infer some general trends. We see that erosion tends to increase stochastic cyclicality \mathcal{K} and reduce alpha (giving a bigger tail), and accretion does the opposite, both quite significantly. Thus simply fitting a power law to a simulation or to empirical data requires some confidence that accretion and erosion processes are not significant in the modelled process. Furthermore, partial shattering tends to leave the exponent little changed but enhances \mathcal{K} . Power-law fragmentation can shift the exponent a little more but still enhances \mathcal{K} . Our results for such power-law fragmentation, where we have a continuous independent variable available to be adjusted, suggest that output should be a continuous function of the model variation, but it is beyond the scope of the present study

to quantify this mapping precisely. Overall, even if the power law seems to fit well, all sorts of non-shattering fragmentation can enhance the \mathcal{K} that we already know exists in the base model: simply put, a well-fitted power law does not indicate that this is necessarily all that is going on.

Gel-shatter cyclicity only occurs in certain narrow regimes in which coalescence can create a gel but size-biased shattering inevitably leads to its dissolution into monomers. Such regimes are, unsurprisingly, less robust. Again, small amounts of accretion and erosion do not have disproportionate effects, but larger amounts can reduce or remove gel-shatter cyclicity, and move the system close to a steady state. The various alternative fragmentation processes (random partition, stick-breaking and CRP) can easily prevent gel-shatter cyclicity: even when a single large cluster forms, it may merely fragment somewhat and re-form in a stochastic but time-symmetric, constant manner. It seems that some degree of shattering into monomers, rather than a more varied distribution of fragments, is necessary to observe true stochastic gel-shatter cycles. A natural path for future analytical work would be to explore and delineate the boundary between a truly steady state and oscillations, perhaps beginning with a linear stability analysis of the evolution equations [44].

Finally, let us return to our opening context, of war and political violence. On the basis of our results, we can say that Richardson's law for the distribution of wars [1] and modern results on deadly events in insurgencies [2] are both consistent with a broad class of CF models, broader than the base model and its generalizations treated deterministically in [7]. Furthermore our findings suggest that the core dynamics are somewhat robust to the behaviour of individuals – captured through the processes of accretion and erosion in Becker-Döring dynamics – providing some explanation of the consistency of the power-law findings in the operations research literature. Nevertheless one should also reasonably expect that an imputed steady state is not the full story, since endogenous (in addition to exogenous) factors can naturally produce not merely stochastic variation but some measure of time-asymmetric cyclicity, even if there is no clear gelation-like phenomenon.

Acknowledgment We should like to thank Stephen Connor, Gustav Delius and Dmitri (Mitya) Pushkin for discussions, and George Constable and Kevin Glazebrook for comments. We also thank the anonymous reviewers for their helpful comments, especially for drawing our attention to the implication that our results hold for a wider class of CF models. This project was undertaken on the Viking Cluster, which is a high performance computer facility provided by the University of York. We are grateful for computational support from the University of York High Performance Computing service, Viking and the Research Computing team.

Author contributions

BTF conducted the simulations. NJM and AJW conceived and supervised the project. All authors participated in writing and editing the article.

Funding Brennen Fagan would like to thank the Department of Mathematics at the University of York for a PhD studentship, and the Leverhulme Centre for Anthropocene Biodiversity at the University of York for a postdoctoral position, Grant Number RC-2018-021. For the purpose of open access a Creative Commons Attribution (CC BY) licence is applied to any Author Accepted Manuscript version arising from this submission.

Data availability statement Code and data to generate the plots are available at <https://github.com/Brennen-Fagan/Robustness-Coal-Frag/>.

Declarations

Conflict of interest The authors have no financial or proprietary interests in any material discussed in this article.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

Appendix A: Mathematical details

A.1 Statistical definitions

In Sect. 2, we employ statistical methodologies adapted from Clauset et al. [24]. Here, we briefly summarize details of the method. We begin by noting maximum likelihood estimation (MLE), in which a given distributional form is fitted by identifying the parameters that yield the distribution that maximizes the probability of generating a given sample conditional on the parameters (termed the likelihood). For a discrete power-law distribution with finite range and random sample of cluster sizes X_1, X_2, \dots, X_n , the estimators for each parameter are

$$\hat{x}_{\min} = \min(X_1, X_2, \dots, X_n) \quad (\text{A1})$$

for the minimum of the range of support,

$$\hat{x}_{\max} = \max(X_1, X_2, \dots, X_n) = k_{\max} \tag{A2}$$

for the maximum of the range of support, and

$$\frac{\sum_{x=x_{\min}}^{x_{\max}} x^{-\hat{\alpha}} \log(x)}{\sum_{x=x_{\min}}^{x_{\max}} x^{-\hat{\alpha}}} = \frac{1}{n} \sum_{i=1}^n \log(X_i) \tag{A3}$$

(numerically solved) for the exponent of the power-law distribution [25].

To construct the Kolmogorov–Smirnov version of MLE (KS-MLE), we do not employ (A1). Instead, we use the Kolmogorov–Smirnov distance between a cumulative distribution function G and the data, defined as

$$\text{KS}(G, X) = \max_x |G(x) - \frac{1}{n} \sum_{i=1}^n \mathbf{1}(X_i \leq x)| \tag{A4}$$

where $|\cdot|$ is the absolute value and $\mathbf{1}(l)$ is 1 if condition l is true and false otherwise. Equation (A4) measures the distance as the largest difference between the cumulative distribution function and the empirical distribution function. The difference between MLE and KS-MLE is that KS-MLE picks x_{\min} to be the x which minimizes (A4). This requires a brute force search over the data [24]. While KS-MLE has the disadvantage of discarding data, it has the advantage that it only fits the power-law tail of a distribution, which can provide less biased estimates if the data is not solely power-law distributed (*e.g.* ones-inflated).

A.2 Characteristic times

Throughout the paper, we have used the cyclicity order parameter \mathcal{K} and the dimensionless number r to relate our modified coalescence and fragmentation systems to the original unmodified system. The parameter \mathcal{K} and number r were used previously by us to describe the observed stochastic gel-shatter cycles when simulating coalescence and fragmentation systems [9].

On the computational timescale, the number r is the ratio of the characteristic time to gelation to the characteristic time to shattering; it captures the balance of whether the system is dominated by the coalescence or the fragmentation process. We measure these two times as

$$T_g = M/\hat{K} \tag{A5}$$

for the characteristic time to gelation and

$$T_f = 1/\hat{F} \tag{A6}$$

for the characteristic time to shattering. (Note that in our simulations mass is conserved, so that analytic definitions of gelation based on the failure of mass conservation [8] are inappropriate.) The time to gelation scales as the typical number of computational steps required to reach the full coalescence of the simulation in the

absence of fragmentation. In contrast, the characteristic time to shattering scales as the typical number of computational steps required to reset the system to its pre-gel state.

If T_g is small and T_f large, then the system achieves gelation quickly and remains in the state for a long period of time, with the mean computational recurrence time scaling (*i.e.* the number of computational steps between successive shatterings of the largest cluster) as T_f . If T_g is large and T_f small, then the system never successfully stays in the gel state for an appreciable length of time. If they are balanced, then with some assumptions the mean recurrence time can be shown to be proportional to $\sqrt{T_g T_f}$ [9]. In this regime, the system experiences a transitory, stochastic, but highly repetitive behaviour in which the largest cluster in the system tends towards gelation only to be quickly fragmented before the sol is too much reduced. This leaves the next largest cluster to repeat the process, which we term stochastic cyclicity following the discussions of stochasticity and cyclicity by McKane and Newman [10].

Finally, it has been shown that the cyclicity order parameter \mathcal{K} , equation (4), responds strongly and non-linearly to r [9]. Hence, \mathcal{K} provides a diagnostic of the behaviour by indicating large changes in the behaviour of the largest cluster. We focus on \mathcal{K} , which more directly tracks the change in the largest cluster, instead of the mean recurrence time because it is not clear if recurrence has been achieved in cases of partial fragmentation where a largest cluster always exists.

A.3 Simulation details

Our simulation details are broken into two halves, due to the nature of the processes we have considered in the main text. In the first half, we consider accretion and erosion which are supplementary processes to the main system dynamics. In the second half, we consider direct modifications to the fragmentation function, such as halving, partitioning, or using the Chinese Restaurant Process.

A.3.1 Becker–Döring dynamics

Our first modification to the processes of coalescence and fragmentation was to implement additional Becker–Döring dynamics in the form of accretion and attrition. Ours differs from standard implementations due to the nature of our time scale (discrete/computational time). We write accretion and erosion rates as a_+ and a_- . Then for accretion, we implement for $k \geq 1$

$$n_k \xrightarrow{ka_+} n_{k+1} - n_1$$

where we abuse notation to indicate mass conservation without a rate-dependence on monomers. For erosion, we have for $k \geq 2$

$$n_k \xrightarrow{ka_-} n_{k-1} + n_1.$$

Expanding (1) with size-biased Becker-Döring dynamics results in

$$\begin{aligned} \frac{dn_k}{dt} = & \frac{1}{2} \sum_{i=1}^{k-1} K(i, (k-i))n_i n_{k-i} - n_k \sum_{i=1}^{\infty} K(i, k)n_i \\ & + \frac{1}{k} \sum_{i=k+1}^{\infty} F(i, k)in_i - n_k \sum_{i=1}^k F(k, i) \\ & + a_+((k-1)n_{k-1} - kn_k) \\ & + a_-((k+1)n_{k+1} - kn_k) \end{aligned} \tag{A7}$$

for $k \geq 2$. For $k = 1$,

$$\begin{aligned} \frac{dn_1}{dt} = & -n_1 \sum_{i=1}^{\infty} K(i, 1)n_i + \sum_{i=2}^{\infty} F(i, 1)in_i \\ & + (a_- - a_+) \sum_{i=1}^{\infty} in_i - a_+n_1 + 2a_-n_2 \end{aligned} \tag{A8}$$

(taking $F(1, 1) = 0$).

However, Eqs. (A7, A8) are specifically written for the case where accretion and erosion can repeatedly target the same cluster in a time step. We noted in the main text that enforcing a uniqueness condition, equivalent to having a cluster become unable to accrete or erode for a short period after having done so, can greatly affect the resulting distribution (Figs. 7 and 11, top-right and bottom-right). While conceptually straightforward, this greatly increases the number of equations, as we implement the unreactive types for accretion and erosion separately (*e.g.* something that has accreted can still be eroded, coalesced, or fragmented). We avoid this problem in our implementation by fixing the rate of accretion and erosion for each computational time step and processing them simultaneously.

A.3.2 Fragmentation functions

Turning our attention now to modifications of the fragmentation function, our modifications all take the form

$$F(i, k) = \hat{F} \frac{i}{M} \sum_{p \in P(i)} \frac{p}{i} \delta_{k,p}, \tag{A9}$$

where \hat{F} is the constant reaction rate of fragmentation with the standard size-bias i and normalized for system size M . The new term, in comparison to (3), is the sum over parts p of the partition of i , $P(i)$, with weighting $\frac{p}{i}$. Whereas the prefactor (outside the sum) in (A9) controls the rate at which clusters are fragmented, the summand describes how each cluster is broken up. Hence $\sum_k \sum_{p \in P(i)} \frac{p}{i} \delta_{k,p} = 1$.

For example, when performing the halving fragmentation (Fig. 8, top-left),

$$F(i, k) = \hat{F} \frac{i}{M} \frac{1}{i} (\lceil i/2 \rceil \delta_{k, \lceil i/2 \rceil} + \lfloor i/2 \rfloor \delta_{k, \lfloor i/2 \rfloor}) \tag{A10}$$

Here, the equivalent number of monomers (“mass”) is split between the number of clusters of size $\lceil i/2 \rceil$ and of size $\lfloor i/2 \rfloor$. Then the factor of i/k in (1) converts the split cluster from its original unit (number of clusters of size i) to its new unit (number of clusters of size k). The i/k and k/i naturally cancel to give one cluster of size $\lceil i/2 \rceil$ and one of size $\lfloor i/2 \rfloor$.

Similarly, when shattering the smaller cluster and permitting other fractions as in the remainder of Fig. 8, we use

$$F(i, k) = \hat{F} \frac{i}{M} \frac{1}{i} (q \delta_{k, \lceil qi \rceil} + (1-q) \delta_{k, 1}), \tag{A11}$$

where $q \in (0, 1)$ is the fraction of the cluster that remains unshattered.

A.3.3 Stochastic fragmentation

Returning to (A9), we now detail the stochastic processes we use to partition a cluster of size i . First, we consider constructing a partition uniformly at random from the space of partitions. We implement algorithm 5 of Arratia and DeSalvo [40].

```

1: procedure PARTITION( $i$ ) ▷  $i \in \mathbb{N} \setminus \{0\}$ 
2:   if  $i == 1$  then
3:     return 1.
4:   end if
5:   while True do
6:     Sample  $X_2, \dots, X_h, \dots, X_i \sim \text{Geometric}(1 - \exp(-h\pi/\sqrt{6i}))$ .
7:     Set  $X_1 = i - \sum_{j=2}^i jX_j$ .
8:     Sample  $U \sim \text{Uniform}_C(0, 1)$ .
9:     if  $X_1 \geq 0, U < \exp(X_1\pi/\sqrt{6i})$  then
10:      return  $X_1, \dots, X_i$ .
11:    end if
12:  end while
13: end procedure

```

(Uniform_C refers to a continuous uniform distribution with the provided bounds.) This algorithm is referred to as a probabilistic divide-and-conquer algorithm, and we direct the reader to Arratia and DeSalvo [40], which relies on work by Fristedt [45], for further details. Note that the Partition procedure returns the number of clusters of a given size (*e.g.* two monomers and a dimer would be returned as (2, 1, 0, 0)), rather than the sizes of clusters (*e.g.* (1, 1, 2)), but converting between the two is straightforward.

Our implementation of discrete stick-breaking (Fig. 9, top-right) is straightforward:

```

1: procedure BREAK( $i$ ) ▷  $i \in \mathbb{N} \setminus \{0\}$ 
2:   Initialise Parts as an empty list.

```

```

3:   Initialise sum = 0.
4:   while sum < i do
5:     Set  $p' \sim \text{Uniform}_D(1, i - \text{sum})$ .
6:     Append  $p'$  to Parts.
7:     Update  $\text{sum} \leftarrow \text{sum} + p'$ .
8:   end while
9:   return Parts.
10: end procedure

```

(Uniform_D refers to a discrete uniform distribution with the provided inclusive bounds.) Note that, in contrast to Partition, Break returns the cluster sizes themselves. While it is not covered in the text, we did briefly explore replacing Uniform_D with a Beta – Binomial distribution with Beta parameters (2, 4), (3, 3), and (4, 2).

We now prove that for all i, j with $i \geq j \geq 1$

$$E[n_j | \text{Break}(i)] = \frac{1}{j}. \tag{A12}$$

Consider first $i = 1$, for which

$$\Pr(\text{Break}(1) = \{1\}) = 1.$$

For $i = 2$, instead we have

$$\Pr(\text{Break}(2) = \{2\}) = \frac{1}{2}$$

due to the draw from the Uniform_D(1, 2) distribution. In the complementary case, a cluster of length 1 is broken off and then the problem reduces to the $i = 1$ case. Hence, $\Pr(\text{Break}(2) = \{1, 1\}) = 1/2$ as well. For $i = 1$ the expectation of the number of clusters of size 1 is

$$E[n_1 | \text{Break}(1)] = 1,$$

and the expectations for n_1 and n_2 for $i = 2$ are

$$E[n_1 | \text{Break}(2)] = 1,$$

and

$$E[n_2 | \text{Break}(2)] = \frac{1}{2}.$$

Writing p'_k for the size of the k th cluster broken, after p'_1 is determined, we saw above that the problem reduces to Break($i - p'_1$) for $i = 1, 2$ weighted by the probability of entering that state. Assume now that the claim holds for all $j \geq 1$ with $j \leq i$, *i.e.*

$$E[n_j | \text{Break}(i)] = \frac{1}{j}.$$

For $i + 1$, we now have trivially that

$$E[n_{i+1} | \text{Break}(i + 1)] = \frac{1}{i + 1}.$$

For arbitrary j now

$$\begin{aligned}
 E[n_j | \text{Break}(i + 1)] &= \frac{1}{i + 1} \sum_{p'_1=1}^{i+1} \delta_{j=p'_1} + E[n_j | \text{Break}(i + 1 - p'_1)] \\
 &= \frac{1}{i + 1} \sum_{p'_1=1}^{i+1} \delta_{j=p'_1} + \frac{1}{j} \delta_{j \leq i+1-p'_1} \\
 &= \frac{1}{i + 1} \left(1 + \frac{1}{j} (i + 1 - j) \right) \\
 &= \frac{1}{j}.
 \end{aligned}$$

Moving on to the Chinese Restaurant Process, the process can be described as a series of monomers **m** coalescing one-by-one with existing clusters of sizes **c** or forming a new (unit size) cluster. With the first monomer assigned its own cluster and n clusters currently, for the j th monomer we write

$$\Pr(m_j = k | \mathbf{m}, \phi, \theta) = \begin{cases} \frac{c_k - \phi}{j - 1 + \theta} & 1 \leq k \leq n \\ \frac{n\phi + \theta}{j - 1 + \theta} & k = n + 1 \end{cases} \tag{A13}$$

where ϕ denotes the penalty to existing tables and θ denotes the attraction to creating a new table [42]. Throughout this paper, we take $\theta = 1$. It is known that the Chinese Restaurant Process produces a power-law distribution on the sizes of the clusters with exponent $1 + \phi$ [41, 42, 46].

For completeness, our implementation can be written as follows.

```

1: procedure CRP( $i, \phi, \theta$ ) ▷
    $i \in \mathbb{N}_{\setminus\{0\}}, 0 \geq \phi < 1, \theta \geq -\phi$ 
2:   Initialise Clusters as an indexed list {1}.
3:   Sample  $p_2 \dots p_i \sim \text{Uniform}_C(0, 1)$ .
4:   Initialise NewCluster  $\leftarrow \theta + \phi$ .
5:   for  $j \leftarrow 2, \dots, i$  do
6:     Set  $\theta \leftarrow \theta + 1$ .
7:     Set  $p \leftarrow p_j \times \theta - \text{NewCluster}$ .
8:     if  $p < 0$  then ▷ Add Cluster.
9:       Append 1 to Clusters.
10:      Update NewCluster  $\leftarrow \text{NewCluster} + \phi$ 
11:     else ▷ Find Cluster.
12:       Initialise  $k \leftarrow 0$ .
13:       while  $p > 0$  do
14:         Update  $k \leftarrow k + 1$ .
15:         Update  $p \leftarrow p - (\text{Clusters}[k] - \phi)$ .
16:       end while
17:       Update Clusters[ $k$ ]  $\leftarrow \text{Clusters}[k] + 1$ .
18:     end if
19:   end for
20:   return Clusters.
21: end procedure

```

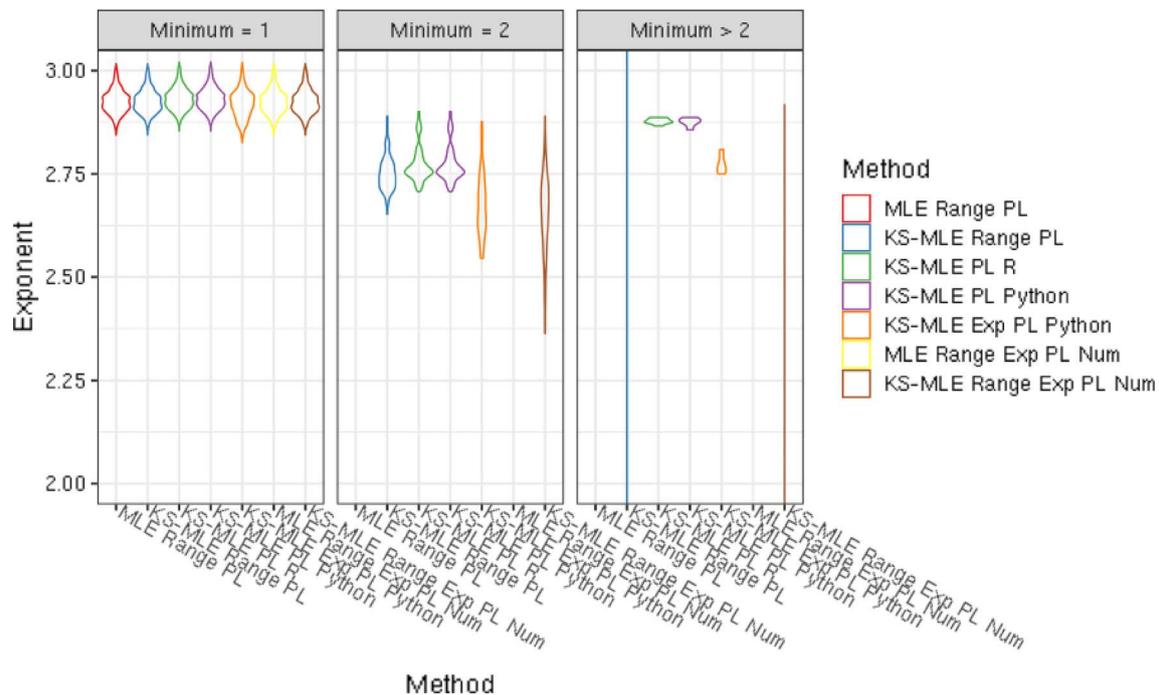


Fig. 15 Distribution of power-law exponents for various fitting methods. Power-laws are fitted to a coalescence and fragmentation process simulated with size-biased kernels as in (2, 3), $M = 10^4$, $\hat{F} = 0.3$, and 20,000 samples. We write MLE and KS-MLE as in the main text, “PL” for power law, “Range” for finite range, “Exp” for an exponential cut-off, “R” for the `powerlaw` package, “Python” for the `powerlaw` package, and “Num” for the numerical maximum likelihood estimation in the `bbmle` package. The theoretical value is 2.5, which all methods deviated from with varying amounts of bias

A.4 Intrinsic estimation bias

In the main text and to match our expectations of fitting power laws in operations research [e.g. 2], we restrict ourselves to fitting discrete power-law distributions with finite ranges via MLE or KS-MLE methods, adapting existing methods in the literature [24]. We show in Sect. 2 that these methods exhibit bias in their estimation of the power-law exponent. This raises the question of whether more distribution-specific methods would better fit the data and, if so, how much better.

As a robust proof-of-concept, we fitted one simulation that should be in steady state (size-biased kernels as in (2, 3), $M = 10^4$, $\hat{F} = 0.3$, and 20,000 samples) with a variety of methods. Using the additional short-hand of “Exp” for exponential cut-off and “Range” for finite range (or finite support), we show in Fig. 15 the fitted power-law exponents using (1) our existing methods, (2) power-law fitting routines in R’s `powerlaw` package [27] and Python’s `powerlaw` package [26], and (3) direct maximum likelihood estimation [47] of a Range-Exp-power-law distribution using MLE and KS-MLE methods. Despite the additional controls and fitting parameters, there is still a large (albeit reduced) bias in the distribution of the power-law exponent numerically fitted via KS-MLE with Range and Exp controls. In this case, when the KS-MLE selects $x_{\min} = 1$ fits are essentially equivalent to the base MLE case and this occurs 79.5% of the time. When $x_{\min} = 2$, 16.1% of the time,

94.4% of fitted power-law exponents are larger than 2.5 (mean 2.66, median 2.67). Visually assessing the other results, we can determine that the KS component is the most capable of reducing bias, followed by the Exp component, and finally the Range component.

References

1. L.F. Richardson, *Statistics of deadly quarrels* (Stevens & Sons, London, 1960)
2. J.C. Bohorquez, S. Gourley, A.R. Dixon, M. Spagat, N.F. Johnson, Common ecology quantifies human insurgency. *Nature* **462**, 911–914 (2009). <https://doi.org/10.1038/nature08631>
3. M. Spagat, N.F. Johnson, S.V. Weezel, Fundamental patterns and predictions of event size distributions in modern wars and terrorist campaigns. *PLOS One* **13**(10), 1–13 (2018). <https://doi.org/10.1371/journal.pone.0204639>
4. M. Spagat, S. Weezel, D.D.J. Restrepo, M. Zheng, N.F. Johnson, Unifying casualty distributions within and across conflicts. *Heliyon* **6**(8), 04808 (2020)
5. R. Zwetsloot, Testing Richardson’s law: a (cautionary) note on power laws in violence data. SSRN article no. 3112804 (2018). <https://doi.org/10.2139/ssrn.3112804>
6. H.A. Simon, On a class of skew distribution functions. *Biometrika* **42**(3/4), 425–440 (1955). <https://doi.org/10.2307/2333389>

7. B. Rusczycki, B. Burnett, Z. Zhao, N.F. Johnson, Relating the microscopic rules in coalescence-fragmentation models to the cluster-size distribution. *Eur. Phys. J. B* **72**, 289–302 (2009)
8. J. Banasiak, W. Lamb, P. Laurençot, Analytic methods for coagulation-fragmentation models. In: Chapman & Hall/CRC Monographs and Research Notes in Mathematics. CRC Press, London, England (2019). <https://doi.org/10.1201/9780429280320>
9. B.T. Fagan, N.J. MacKay, D.O. Pushkin, A.J. Wood, Stochastic gel-shatter cycles in coalescence-fragmentation models. *Europhys. Lett.* (2021)
10. A.J. McKane, T.J. Newman, Predator-prey cycles from resonant amplification of demographic stochasticity. *Phys. Rev. Lett.* **94**(21), 218102 (2005). <https://doi.org/10.1103/PhysRevLett.94.218102>
11. H. Tanaka, S. Inaba, K. Nakazawa, Steady-state size distribution for the self-similar collision cascade. *Icarus* **123**(2), 450–455 (1996). <https://doi.org/10.1006/icar.1996.0170>
12. T. Birnstiel, C.W. Ormel, C.P. Dullemond, Dust size distributions in coagulation/fragmentation equilibrium: numerical solutions and analytical fits. *Astron. Astrophys.* **525**, 11 (2011). <https://doi.org/10.1051/0004-6361/201015228>
13. D.J. Aldous, Deterministic and stochastic models for coalescence (aggregation and coagulation): a review of the mean-field theory for probabilists. *Bernoulli* **5**(1), 3–48 (1999)
14. D.O. Pushkin, H. Aref, Bank mergers as scale-free coagulation. *Phys. A: Stat. Mech. Appl.* **336**(3), 571–584 (2004). <https://doi.org/10.1016/j.physa.2003.12.056>
15. S. Datta, G.W. Delius, R. Law, M.J. Plank, A stability analysis of the power-law steady state of marine size spectra. *J. Math. Biol.* **63**(4), 779–799 (2011). <https://doi.org/10.1007/s00285-010-0387-z>
16. R. D’Hulst, G.J. Rodgers, Exact solution of a model for crowding and information transmission in financial markets. *Int. J. Theor. Appl. Fin.* **03**(04), 609–616 (2000). <https://doi.org/10.1142/S0219024900000784>
17. R.C. Ball, C. Connaughton, P.P. Jones, R. Rajesh, O. Zaboronski, Collective oscillations in irreversible coagulation driven by monomer inputs and large-cluster outputs. *Phys. Rev. Lett.* **109**(16) (2012). <https://doi.org/10.1103/physrevlett.109.168304>
18. R.L. Pego, J.J.L. Velázquez, Temporal oscillations in Becker–Döring equations with atomization. *Nonlinearity* **33**(4), 1812–1846 (2020). <https://doi.org/10.1088/1361-6544/ab6815>
19. S.A. Matveev, P.L. Krapivsky, A.P. Smirnov, E.E. Tyrtshnikov, N.V. Brilliantov, Oscillations in aggregation-shattering processes. *Phys. Rev. Lett.* **119**, 260601 (2017). <https://doi.org/10.1103/PhysRevLett.119.260601>
20. A. Clauset, F.W. Wiegel, A generalized aggregation-disintegration model for the frequency of severe terrorist attacks. *J. Conflict Resolut.* **54**(1), 179–197 (2010). <https://doi.org/10.1177/0022002709352452>
21. A.E. Kyrianiou, S.W. Pagett, T. Rogers, Universality in a class of fragmentation-coalescence processes. *Annales de l’Institut Henri Poincaré, Probabilités et Statistiques* **54**(2), 1134–1151 (2018). Institut Henri Poincaré
22. R.J. Field, R.M. Noyes, Oscillations in chemical systems. iv. limit cycle behavior in a model of a real chemical reaction. *J. Chem. Phys.* **60**(5), 1877–1884 (1974)
23. K. Tkacova, A. Idler, N. Johnson, E. López, Explaining conflict violence in terms of conflict actor dynamics. *Sci. Rep.* **13**(1), 21187 (2023). <https://doi.org/10.1038/s41598-023-48218-x>
24. A. Clauset, C.R. Shalizi, M.E.J. Newman, Power-law distributions in empirical data. *SIAM Rev.* **51**(4), 661–703 (2009)
25. H. Zhu, Y. Xie, M. Xu, Discrete truncated power-law distributions. *Aust. New Zeal. J. Stat.* **58**(2), 197–209 (2016). <https://doi.org/10.1111/anzs.12162>
26. J. Alstott, E. Bullmore, D. Plenz, powerlaw: a python package for analysis of heavy-tailed distributions. *PLOS One* **9**(1), 1–11 (2014). <https://doi.org/10.1371/journal.pone.0085777>
27. C.S. Gillespie, Fitting heavy tailed distributions: the poweRlaw package. *J. Stat. Software* **64**(2), 1–16 (2015)
28. H. Wickham, *Ggplot2: Elegant Graphics for Data Analysis*. Springer-Verlag, New York (2016). <https://ggplot2.tidyverse.org>
29. M. Smoluchowski, Versuch einer mathematischen theorie der koagulationskinetik kolloider lösungen. *Zeitschrift fuer physikalische Chemie* **92**, 129–168 (1916)
30. P. Erdős, A. Rényi, On the evolution of random graphs. *Publ. Math. Inst. Hung. Acad. Sci* **5**(1), 17–60 (1960)
31. J.A.D. Wattis, An introduction to mathematical models of coagulation-fragmentation processes: a discrete deterministic mean-field approach. *Physica D: Nonlinear Phenomena* **222**(1–2), 1–20 (2006)
32. J.M. Ball, J. Carr, O. Penrose, The Becker–Döring cluster equations: basic properties and asymptotic behaviour of solutions. *Commun. Math. Phys.* **104**, 657–692 (1986). <https://doi.org/10.1007/BF01211070>
33. D.O. Pushkin, H. Aref, Self-similarity theory of stationary coagulation. *Phys. Fluids* **14**(2), 694–703 (2002). <https://doi.org/10.1063/1.1430440>
34. N. Brilliantov, P.L. Krapivsky, A. Bodrova, F. Spahn, H. Hayakawa, V. Stadnichuk, J. Schmidt, Size distribution of particles in Saturn’s rings from aggregation and fragmentation. *Proc. Natil. Acad. Sci.* **112**(31), 9536–9541 (2015). <https://doi.org/10.1073/pnas.1503957112>. <http://www.pnas.org/content/112/31/9536.full.pdf>
35. N.F. Johnson, M. Zheng, Y. Vorobyeva, A. Gabriel, H. Qi, N. Velasquez, P. Manrique, D. Johnson, E. Restrepo, C. Song, S. Wuchty, New online ecology of adversarial aggregates: ISIS and beyond. *Science* **352**(6292), 1459–1463 (2016). <https://doi.org/10.1126/science.aaf0675>
36. F.P. Costa, Existence and uniqueness of density conserving solutions to the coagulation-fragmentation equations with strong fragmentation. *J. Math. Anal. Appl.* **192**(3), 892–914 (1995). <https://doi.org/10.1006/jmaa.1995.1210>
37. R. Becker, W. Döring, Kinetische behandlung der keimbildung in übersättigten dämpfen. *Annalen der Physik* **416**(8), 719–752 (1935). <https://doi.org/10.1002/andp.19354160806>
38. N.V. Brilliantov, W. Otieno, P.L. Krapivsky, Nonextensive supercluster states in aggregation with fragmentation. *Phys. Rev. Lett.* **127**(25), 250602 (2021)
39. W. Otieno, N.V. Brilliantov, P.L. Krapivsky, Supercluster states and phase transitions in aggregation-

- fragmentation processes. *Phys. Rev. E* **108**(4), 044142 (2023)
40. R. Arratia, S. DeSalvo, Probabilistic divide-and-conquer: a new exact simulation method, with integer partitions as an example. *Combin. Prob. Comput.* **25**(3), 324–351 (2016). <https://doi.org/10.1017/S0963548315000358>. [arXiv:1110.3856](https://arxiv.org/abs/1110.3856) [math.PR]
 41. J. Pitman, *Combinatorial Stochastic Processes*, Lecture Notes in Mathematics, vol. 1875. (Springer, Berlin, 2006)
 42. S. Goldwater, T.L. Griffiths, M. Johnson, Producing power-law distributions and damping word frequencies with two-stage language models. *J. Mach. Learn. Res.* **12**, 2335–2382 (2011)
 43. M.E.J. Newman, *Networks: an introduction* (Oxford University Press, Oxford, 2010)
 44. N.V. Brilliantov, W. Otieno, S.A. Matveev, A.P. Smirnov, E.E. Tyrtshnikov, P.L. Krapivsky, Steady oscillations in aggregation-fragmentation processes. *Phys. Rev. E* **98**, 012109 (2018)
 45. B. Fristedt, The structure of random partitions of large integers. *Trans. Am. Math. Soc.* **337**(2), 703–735 (1993)
 46. B. Bassetti, M. Zarei, M. Cosentino Lagomarsino, G. Bianconi, Statistical mechanics of the “chinese restaurant process”: lack of self-averaging, anomalous finite-size effects, and condensation. *Phys. Rev. E* **80**(6), 066118 (2009). <https://doi.org/10.1103/PhysRevE.80.066118>
 47. B. Bolker, R Development core team: *bbmle: tools for general maximum likelihood estimation* (2020). R package version 1.0.23.1. <https://CRAN.R-project.org/package==bbmle>