POLITECNICO DI TORINO

MASTER's Degree in COMPUTER ENGINEERING

MASTER's Degree Thesis

Popularity Dynamics on Bluesky: A Large-Scale Stochastic Analysis of High-Volatility Environments

Supervisors	Candidate
Prof. Emilio LEONARDI	Giovanni STINÀ
Prof. Franco GALANTE	

March 2026

Thesis Title

Popularity Dynamics on Bluesky: A Large-Scale Stochastic Analysis of High-Volatility Environments

Thesis Author

Giovanni Stinà

Abstract

This thesis investigates the stochastic dynamics of popularity on Bluesky, a decentralized online social network, to challenge the idea that follower count defines influence. By analyzing the ecosystem, we quantify the continuous activity required to balance attention decay. We developed a custom crawler to reconstruct full user timelines from the public stream, compiling a longitudinal dataset of over 31 million interactions. We analyzed these dynamics using a Stochastic Mean-Field framework, which models popularity not as a static metric, but as a stochastic process governed by the opposing forces of content amplification and temporal decay. In this study, we leverage a recently developed stochastic model to characterize the platform's governing dynamics. We quantified the structural 'friction', representing the decay of collective attention, and the viral 'volatility' driving popularity growth. This process isolated the specific coefficients governing the system's memory and long-term equilibrium, ultimately revealing an ecosystem defined by extremely short-lived attention.

Finally, we applied this framework to investigate the structural stratification of the online social network. We contrasted the behavioral patterns of established 'Influencer' elite, defined by high accumulated engagement, against the aspiring 'Wannabe' population. This analysis provides a quantitative basis to evaluate the maintenance cost imposed by the platform's rapid decay on emerging actors. Ultimately, we investigate whether the ecosystem's high volatility allows for social mobility or if it tends to reinforce the dominance of established elites.

Table of Contents

List of Figures

List of Tables

Acronyms

API	Application Programming Interface.
AT	Authenticated Transfer.
CCDF	Complementary Cumulative Distribution Function.
c-TF-IDF	Class-based Term Frequency-Inverse Document Frequency.
DID	Decentralized Identifier.
DNS	Domain Name System.
HDBSCAN	Hierarchical Density-Based Spatial Clustering of Applications with Noise.
HIP	Hawkes Intensity Process.
I/O	Input/Output.
JSON	JavaScript Object Notation.
JSONL	JSON Lines.
KS	Kolmogorov-Smirnov.
MLE	Maximum Likelihood Estimation.
NaN	Not a Number.
NLP	Natural Language Processing.
PDS	Personal Data Servers.
Q4	Fourth Quarter.
SDE	Stochastic Differential Equation.
SDK	Software Development Kit.

UMAP	Uniform Manifold Approximation and Projection.
URI	Uniform Resource Identifier.
UTC	Coordinated Universal Time.

Chapter 1

Introduction

1.1 Background and Problem Statement

In the contemporary attention economy, collective attention is a strictly limited resource governed by strong competition and the rapid decay of novelty. Traditionally, centralized social media platforms have managed this scarcity through opaque recommendation algorithms designed for forced engagement. In response, the landscape has experienced a paradigm shift towards decentralized architectures, with Bluesky emerging as a central player [1]. Built upon the Authenticated Transfer (AT) Protocol, Bluesky structurally decouples data hosting, indexing, and moderation services. User records are managed through independent Personal Data Servers (PDS), while Relays aggregate these records into a real-time public stream known as the firehose. This open infrastructure replaces corporate algorithms with independent Feed Generators, offering users true "algorithmic choice" [2].

However, while the technical architecture of Bluesky democratizes content curation, the ecosystem remains bound by the fundamental limits of human cognitive capacity. In such a high-volatility environment, traditional static metrics, such as accumulated follower counts, are inadequate to define true influence . Instead, popularity must be understood as a highly dynamic and diffusive stochastic process. It is continuously shaped by two opposing forces: the viral 'volatility' that drives content amplification, and the structural 'friction' representing the inevitable decay of collective attention [3].

The core problem this thesis addresses is whether the unprecedented technical freedom of decentralized networks translates into actual social mobility, or if it inadvertently reproduces systemic asymmetries. The high volatility and rapid attention decay characteristic of Bluesky impose a continuous maintenance cost on users. It is imperative to investigate how this cost impacts the structural stratification of the network, specifically by contrasting the behavioral patterns of the established 'Influencer' elite against the aspiring 'Wannabe' population. Ultimately, we investigate whether this algorithmic freedom actually helps new users gain visibility, or if it simply keeps the established elites at the top.

1.2 Objectives and Originality

The primary objective of this thesis is to empirically quantify the continuous activity required to balance attention decay on Bluesky, and to investigate whether the platform's high volatility encourages social mobility or reinforces the dominance of established elites.

The originality of this work lies fundamentally in both its empirical and theoretical contributions. On the empirical front, we constructed a massive and highly detailed longitudinal dataset, which is a rare asset in the study of decentralized networks. By developing a custom multi-threaded crawler that interfaces directly with the AT Protocol's firehose, we reconstructed the full 180-day historical timelines of a stratified sample of nearly 30,000 active users. This resulted in a proprietary dataset comprising over 31 million distinct post and repost. This large dataset separates original posts from reposts in order to allows us to directly observe the fast-paced attention economy.

Theoretically, we move beyond static network analyses by leveraging this dataset within a Stochastic Mean-Field framework. Applying the methodology proposed by Galante et al. [3], we utilize Maximum Likelihood Estimation (MLE) via Grid Search to isolate the specific coefficients governing the system's memory and long-term equilibrium. This combined empirical and theoretical approach provides a robust quantitative basis to evaluate the maintenance cost imposed by the platform on emerging actors.

1.3 Methodology

This research uses a step-by-step methodology to analyze the Bluesky ecosystem. First, we develope a custom multi-threaded crawler. This tool connects directly to the AT Protocol firehose. It extracts public data without relying on restricted corporate APIs. We reconstruct the full 180-day historical timelines for a stratified sample of about 30,000 active users. This extraction built a longitudinal dataset of over 31 million distinct interactions. Within this dataset, we strictly separate original content creation from network propagation (repost).

Second, we analyze these data using a Stochastic Mean-Field framework. We model popularity not as a static number, but as a continuous stochastic process. We establish that two opposing forces govern this process: viral volatility, which drives content amplification, and structural friction, which represents the natural decay of collective attention. We characteriz these governing dynamics by applying a recently developed mathematical model presented in [3].

Third, we calibrate this mathematical model using our real-world data. We apply Maximum Likelihood Estimation (MLE) through a Grid Search. This optimization process isolates the specific coefficients that control the system's memory and long-term equilibrium.

Finally, we apply this calibrated framework to study the network's social structure.

We group the users based on their average engagement using a K-Means clustering algorithm. This algorithm divides the population into distinct tiers. It strictly defines the aspiring 'Wannabe' population as the medium engagement class and the established 'Influencer' elite as the high engagement class. We then compare the behavioral patterns of these two specific groups. This quantitative comparison allows us to evaluate the exact maintenance cost that the platform's rapid attention decay imposes on emerging actors.

1.4 Thesis Outline

This thesis contains five main chapters and a conclusion.

- **Chapter One - Introduction** introduces the context of the attention economy. It defines the core research problem and states the main objectives.
- **Chapter Two - Related work** reviews the related literature. First, it explores the technical paradigm shift introduced by Bluesky and the AT Protocol. It details user migration dynamics and network bootstrapping through tools like Starter Packs. Second, it examines the modern attention economy. It explains the limits of cognitive capacity, the rapid decay of novelty, and the bursty nature of online interactions. Third, the chapter reviews the mathematical modeling of popularity. It introduces the Hawkes Intensity Process and Stochastic Differential Equations (SDEs) to separate endogenous virality from exogenous shocks. Finally, it analyzes social stratification. It explains how homophily and algorithmic choices reinforce structural biases, creating an algorithmic glass ceiling that restricts actual social mobility.
- **Chapter Three - Data and Methodology** details the data acquisition and methodology. First, we explain the hybrid pipeline developed to extract the custom feed ecosystem. We call this approach hybrid because it combines web scraping a third-party directory to discover feed URLs with official AT Protocol API calls to retrieve their actual metadata. Next, we describe the four-step sampling strategy for user interactions. We monitored the real-time firehose to identify active users. We performed a short-term activity crawl. We applied K-Means clustering to stratify users into engagement tiers. Finally, we executed a deep historical crawl to capture 180 days of activity. Second, we characterize the resulting datasets. The Feed Landscape Dataset maps over 79,000 custom algorithms. The Deep User Interaction Dataset compiles over 31 million actions from a sample of nearly 30,000 users. In this section, we also explain how we removed saturated timelines to prevent bias in our mathematical models. Finally, we outline the data processing pipelines. We specify the Natural Language Processing (NLP) techniques used to cluster feed descriptions. We also define how we transformed the raw interactions into continuous mathematical metrics. This final step prepares the empirical data for the stochastic models applied in the subsequent chapters.

- **Chapter Four - Bluesky - Exploratory Data Analysis** presents the exploratory data analysis. First, it investigates the custom feed ecosystem. We show that feed creation occurs in sudden bursts driven by external platform migrations. We observe a power law in creator productivity and expose a "quantity versus quality" paradox. We define the "Zero-Like Phenomenon," demonstrating that the majority of algorithms receive zero engagement. We also prove that pre-existing social capital creates a glass ceiling for feed success, while longevity provides no advantage. Second, it examines the user interaction dataset. We stratify the population into Low, Medium, and High engagement tiers using a K-Means clustering algorithm. We analyze the distribution of user activity and explain the temporal bias caused by strict API data limits. Finally, we build a multi-dimensional heuristic classifier to separate automated bots from legitimate human users within the saturated class. We analyze the mathematical growth regimes of both classes to evaluate their structural stability.

- **Chapter Five - Global Dynamics of Popularity and Social Stratification** models the global dynamics of popularity and social stratification. First, we introduce the Stochastic Mean-Field framework. We model popularity as a continuous stochastic process driven by physiological decay, endogenous jumps, and exogenous shocks. We also define the conditions for system ergodicity and the interpretation of structural parameters. Second, we calibrate the model on the empirical data. We apply Maximum Likelihood Estimation through a Grid Search, adopting the methodology from [3]. We estimate the global parameters to quantify the system's viral capacity and memory. We then validate the model by testing the statistical distribution of the popularity jumps. Finally, we analyze social stratification by comparing the Influencer and Wannabe regimes. We perform a class-specific calibration to measure quality variance and temporal burstiness. We demonstrate that Influencers operate in a near-critical growth regime, while Wannabes remain trapped in a sub-critical state.

- **Chapter Six - Conclusion** summarizes the empirical findings. It discusses the actual level of social mobility within the decentralized network.

Chapter 2

Related Work

2.1 The Paradigm Shift: From Centralized Algorithms to "Algorithmic Choice"

In recent years, the social media landscape has experienced a new transformation, moving away from centralized models to explore decentralized architectures. Within this context, Bluesky has emerged as a central player, reaching 30 million users at the end of January 2025 [4]. The platform represents a successful experiment in creating a decentralized "digital town square" that provides a seamless user experience (speed, global search, simple onboarding) while maintaining an open technical architecture that is resistant to the dominance of a single entity.

2.1.1 The AT Protocol Architecture and Identity Portability

From a technical perspective, Bluesky operates as a decentralized social network layered atop the AT Protocol (Authenticated Transfer), an open framework designed for large-scale social applications [1]. Unlike centralized systems, the protocol architecture is characterized by a structural decoupling of data hosting, indexing, and moderation services. This separation is intended to provide users with a "credible exit". The technical agency to migrate data across different providers without losing their social graph or identity [1].

The system achieves this by decomposing the infrastructure into cooperative services. User records are managed through Personal Data Servers (PDS), which function as independent, silent, and interchangeable repositories [2, 5]. Aggregation occurs through Relays, which monitor these PDSs to generate a real-time stream of all public records known as the *firehose* [1]. This stream is subsequently consumed by App Views, which build complex indexes, as reply threads and engagement metrics, to serve the data to the clients.

This infrastructure effectively prevents platform lock-in by utilizing immutable Decentralized Identifiers (DIDs) and DNS-based handles, allowing for identity portability that exceeds the capabilities of contemporary decentralized alternatives.

2.1.2 Migration Dynamics: The Academic Exodus and External Shocks

The transition towards Bluesky has been heavily catalyzed by a mass exodus from centralized platforms like X (formerly Twitter) [6, 7, 4]. A detailed study of the migration of 300,000 scholars has provided crucial insights into the dynamics of this shift [6].

The data reveals significant disciplinary heterogeneity: while transition rates reach 31.3% for Arts and Humanities and 26.8% for Social Sciences, they stop at 13.3% for Medicine and Health, as well as Physical Sciences (14.6%) and Engineering/Computer Science (15.7%). The primary drivers of this migration are not related to traditional academic metrics (such as the h-index or citations), but rather to political expression and network metrics, such as PageRank centrality in the original social graph. Furthermore, influence is asymmetric: users are influenced by their information sources (followees) ten times more than by their own audience (followers).

The literature has identified four main contagion mechanisms in this migration [6]:

1. **Simple Contagion (66.7%):** The primary driving force, where every single contact who migrates increases the probability that the user will follow.

2. **External Shocks (16.2%):** Macroscopic events, such as the 2024 US elections or the ban of X in Brazil, which cause sudden spikes. Politically inactive users have shown an 83% higher propensity to move precisely during these events.

3. **Spontaneous Adoption (12.6%):** Movements by early adopters that are not driven by obvious network stimuli.

4. **Complex Contagion (4.5%):** Adoptions that require multiple and simultaneous social reinforcements.

2.1.3 Network Bootstrapping via Starter Packs

To facilitate the rapid reconstruction of social graphs disrupted by migration, Bluesky introduced *Starter Packs* in June [8]. These tools allow new users to follow curated lists containing up to 150 accounts and 3 algorithmic feeds with a single click.

The impact of these tools on network topology has been systemic. As of January 1, 2025, the community had generated over 335,000 Starter Packs, which, at their peak, processed 43% of daily "follow" operations and generated 19.95% of all new network connections [8]. The effectiveness is also measurable at the individual level: users included in these lists benefit from an 85% increase in followers, receive 70% more likes, and increase their posting activity by 60% [8].

Although the dominant themes reflect the needs of professional migrations (journalism, pro-democracy activism, academic and scientific networks), the literature also points out critical risks. Specifically, Starter Packs have raised concerns regarding

non-consensual inclusion, [8] the creation of target lists for harassment, and the emergence of black markets where creators demand payment for inclusion.

2.1.4 Algorithmic Choice and the Literature Gap

The culmination of the paradigm shift offered by Bluesky lies in its departure from centralized and opaque moderation models. By relying on independent third-party entities (*Labelers* for moderation and *Feed Generators* for algorithms), the platform aims to provide true "algorithmic choice". Users are no longer restricted to a single corporate algorithm designed for forced engagement; instead, they can choose from thousands of *Custom Feeds* developed by the community [1].

However, while the literature has extensively mapped the network topology, migration triggers, and onboarding tools, the actual behavior of this algorithmic marketplace remains largely unexplored. If users are free to choose their own content filters, how is collective attention genuinely distributed? Does the abundance of algorithmic supply guarantee fair visibility, or does it reproduce systemic asymmetries?

It is precisely to address these questions and fill this quantitative void that Chapter 4 of this thesis proposes an empirical investigation into the Custom Feed ecosystem, providing a tangible measurement of the dynamics of attention concentration and scarcity.

2.2 The Attention Economy and the Acceleration of Collective Cycles

The analysis of modern media consumption shows a clear transformation in how society manages attention. Research indicates that collective attention is a limited resource. It is governed by strong competition and by the decay of novelty, two dynamics that are intensifying rapidly.

2.2.1 Limited Capacity and Competitive Dynamics

A key finding in the literature is that individual attention capacity is constant. Even though the diversity of information continues to grow, users cannot process more data than before. Metrics based on Shannon entropy confirm this cognitive limit [9]. This bottleneck forces information units, often called "memes", to compete for survival. Therefore, the visibility of a topic does not depend only on its quality. Instead, it depends on its ability to displace other content within a limited cognitive space [9].

The structural friction that drives this popularity decay is deeply rooted in the economics of attention. Recent sociological frameworks conceptualize attention as a severely limited biological resource that acts as the primary currency of social media [10]. In this context, it is crucial to distinguish between "calcified attention," which represents the static accumulation of past interactions such as follower counts, and "flow attention," which represents the active, continuous engagement required to

sustain visibility . This dual-stream model supports the premise that static metrics alone cannot define influence in high-volatility environments; instead, continuous activity is mandatory to balance the natural dissipation of collective attention [10].

2.2.2 Novelty Decay and the Stretched-Exponential Law

The lifecycle of online content is driven by the decay of novelty. Empirical data shows that attention towards new items does not fade in a linear way. Instead, it follows a "stretched-exponential" law [11].

This distribution suggests there is a natural time limit for interest. After this period, attention inevitably disappears due to habituation. The "novelty factor" [11] is the main driver of initial attention, but its effect is temporary.

2.2.3 The Acceleration of Collective Attention

The most significant recent phenomenon is the acceleration of popularity cycles [12]. Longitudinal analysis shows that popularity gradients are becoming steeper. Topics reach their peak visibility faster, but they are abandoned just as quickly [13]. This happens because the rates of content production and consumption are increasing. Consequently, the time available for any single topic to remain at the center of the debate is compressed.

The abundance of information, combined with users' limits, leads to greater temporal fragmentation [13]. This is evident in social media, but also in slower domains like literature and cinema, suggesting a broad cultural change. However, knowledge-based systems like Wikipedia or scientific citations are an exception. They show much less acceleration, suggesting they follow different mechanisms compared to simple popularity or novelty dynamics.

2.2.4 Temporal Patterns: Burstiness and Priority Queues

The analysis of contemporary digital systems reveals that online popularity and human activity dynamics do not follow random or linear models. While traditional models assume a Poisson distribution implying regularity and predictability, real-world data from web traffic and electronic communications exhibit "fat-tailed" distributions characterized by critical phenomena defined as "bursty" [14]. In this regime, long periods of inactivity are interrupted by sudden and intense spikes of activity.

These dynamics are primarily fueled by priority-based queuing processes. Individuals do not execute tasks randomly; instead, they prioritize specific activities, creating extremely heterogeneous waiting times [14]. Consequently, online success is not merely a gradual accumulation process (the "rich-get-richer" mechanism) but is subject to "rank-shifts" caused by exogenous shocks that catapult specific content to collective attention. This implies that resource management systems based on Poisson assumptions (such as server allocation or traffic handling) are inadequate, as they fail to account for the long wait times generated by these priority queues [14].

2.2.5 Structural Bias and the Algorithmic Glass Ceiling

The introduction of social recommendation algorithms intensifies pre-existing structural biases, creating a phenomenon defined as the "Algorithmic Glass Ceiling" [15]. These systems often appear neutral, but in reality, they increase inequality for specific demographic groups. The main driver is "homophily", which is the tendency of users to connect with people similar to themselves.

The synthesis of data suggests that algorithmic neutrality is illusory when the underlying structure is inherently unfair. The recommendation engine rewards the existing connection structure, thereby reinforcing historical biases [15]. This creates a feedback loop that allows a minority (or a group with more cohesive network behaviors) to dominate the peaks of visibility. As explicitly stated in the literature, this algorithmic glass ceiling "exhibits all the properties of the metaphorical social barrier that prevents groups [...] from reaching equal representation".

Therefore, since recommendation algorithms accelerate the creation of inequality, it is necessary to develop "rejection" or correction methods in algorithmic design that account for homophily parameters to ensure equal visibility opportunities [15].

2.3 Mathematical Modeling of Popularity Dynamics

The analysis of online popularity has evolved from simple observation of linear growth mechanisms ("rich-get-richer") to advanced mathematical models capable of explaining the intermittent and bursty nature of attention. Recent literature defines popularity dynamics as a "critical system", physically comparable to earthquakes or avalanches, characterized by fat-tailed distributions both in the magnitude of events and in the inter-arrival times between them.

2.3.1 Endogenous vs. Exogenous Drivers: The Hawkes Intensity Process (HIP)

Online popularity exhibits the dynamics of a critical system. Ratkiewicz et al. established this by analyzing massive web datasets [12]. Their model separates popularity into two distinct forces: the first one is an endogenous mechanism driven by preferential attachment and the second one consists of random exogenous shocks triggered by external factors. This division explains the fat-tailed distributions of interaction events. The Hawkes Intensity Process (HIP) directly implements this conceptual dichotomy.

The intensity of attention $\lambda(t)$ is defined as a self-exciting point process [16]:

$$\lambda(t) = \mu(t) + \sum_{t_i<t} \phi(t - t_i) \tag{2.1}$$

where $\mu(t)$ represents the exogenous stimuli (e.g., external promotions or news events), while the kernel $\phi(t)$ captures the endogenous contagion, representing the influence of past events on the current probability of new interactions. This distinction is

critical: it allows researchers to separate content that is genuinely "viral" (driven by word-of-mouth) from content that is merely "promoted" by external factors.

While the standard HIP framework captures the temporal evolution of these events, recent advancements in continuous-time network modeling further validate the need to segment the users generating them by activity levels. Multivariate Hawkes process models demonstrate that extreme degree heterogeneity heavily skews temporal network dynamics [17]. Isolating high-activity "influencer" nodes from standard user communities systematically improves the predictive accuracy of these stochastic models [17].

2.3.2 The Mean-Field Approach and Stochastic Differential Equations

While point processes model discrete events, understanding the long-term trajectory of influencers requires a *Mean-Field* approach [3]. In this framework, users do not interact with every single peer but with the "average state" of the system.

The dynamics of influencer popularity $x_i(t)$ are modeled through Stochastic Differential Equations (SDEs) that account for three main components:

- **Drift (α):** The deterministic growth trend.
- **Volatility (σ):** The stochastic noise inherent in the system.
- **Decay (γ):** The restoring force that pushes popularity back towards zero over time.

This formulation allows for the analysis of "Rank Dynamics" as a zero-sum game, where the rise of one node necessitates the decline of others within a finite attention economy.

2.3.3 Empirical Calibration and the "Inertia" Factor

Calibrating these models on real-world data (specifically from Facebook) has revealed crucial insights into the stability of established social networks. First, static metrics such as follower counts have proven inadequate for modeling dynamics; instantaneous engagement metrics are required to capture the true state of the system [3].

Using Maximum Likelihood Estimation (MLE), the literature has identified standard parameters for an other platforms (Facebook):

- **Non-linear Peak Dependence (0.7):** The relationship between an influencer's current status and their ability to generate new peaks is sub-linear.
- **System Inertia ($\tau \approx 128$ days):** This is a critical finding. The "memory" of popularity lasts approximately 128 days. This high inertia implies a slow decay rate.

Furthermore, a positive correlation has been observed between popularity and posting frequency: top influencers publish more often to maintain visibility, although excessive frequency can trigger negative feedback mechanisms (audience fatigue).

2.3.4 System Regimes: Dominance vs. Fair Play

This mathematical model is particularly useful for understanding the state of a social network. The behavior of the system depends on the balance between the decay parameter and the amplification factors:

- **Dominance:** If popularity decays too slowly, top influencers keep all the attention. This creates a closed system where it is very hard to change the rankings.
- **Fair Play:** In this state, new users can emerge if their content is valid ("merit"). This balances out the tendency of algorithms to only promote already famous accounts.

Finally, the platform's specific algorithm (like EdgeRank) determines where the network stands between these two extremes.

2.4 Social Stratification and the Glass Ceiling

The analysis of the Bluesky ecosystem reveals a complex social structure. While the platform offers decentralized transparency, it also reflects deep-seated tensions found in traditional social media. The data highlights a divide between the theoretical promise of an open network and the reality of a stratified environment driven by specific algorithmic dynamics.

2.4.1 The Matthew Effect and Cumulative Advantage

The phenomenon of self-reinforcing inequality, where initial advantage systematically begets further advantage, is broadly defined across natural and social sciences as the Matthew effect [18]. In the context of network science and online social platforms, this principle is formally modeled as preferential attachment. This mechanism dictates that new connections or interactions are not distributed uniformly; instead, they are disproportionately acquired by nodes that are already highly connected or visible.

Extensive empirical evidence demonstrates that this dynamic of cumulative advantage governs the growth of various socio-technical systems, driving the emergence of scale-free topologies and power-law distributions in human attention [18]. In high-volatility environments, preferential attachment implies that visibility is heavily modulated by a user's pre-existing network status rather than just the intrinsic merit of new content. This creates a structural baseline where social mobility is inherently constrained by the system's tendency to continuously amplify established actors [18].

2.4.2 Reactive Growth and Ideological Homogeneity

The expansion of Bluesky has not been linear. Instead, it is characterized as "Reactive Growth". The increase in the user base is strictly correlated with critical events occurring on X (formerly Twitter), such as service interruptions, the introduction of

fees, or the US elections. Data shows that 13.5 million new profiles were registered just between December 2024 and May 2025 [5].

Recent empirical models validate this reactive dynamic. Large-scale data shows that exogenous shocks, such as the US elections or service bans, act as the primary catalyst for mass platform migration. Furthermore, migration probability is significantly higher for users expressing progressive political views, directly reinforcing the ideological homogeneity of the new user base [6].

Unlike generalist platforms, the resulting community shows strong ideological homogeneity, with a marked center-left orientation. This structural alignment directly impacts information quality. A consequence of this homogeneity is the quality of information: the spread of fake news or links from questionable sources is almost non-existent, recorded at only 0.14% of shared links [2].

This homogeneity produces asymmetric polarization. Despite broad political alignment, deep structural divisions persist on specific geopolitical issues, such as the Israel-Palestine or Russia-Ukraine conflicts. These divisions are fundamentally unbalanced. A dominant majority systematically overwhelms small minorities, fostering an ideal environment for echo chambers [5].

2.4.3 Network Topology: The 'Rich-Get-Richer' Elite

This topology is rooted in the foundational scale-free network model [19]. Barabási and Albert first demonstrated that complex networks grow through a process called preferential attachment. In this process, highly connected nodes attract new links much faster than smaller nodes. This exact mechanism creates the heavy-tailed distributions seen in modern social platforms, a dynamic extensively formalized in subsequent statistical mechanics reviews [20].

Large-scale empirical data defines the structure of this growing population. Network analysis of over 36 million accounts demonstrates highly heterogeneous, heavy-tailed follower distributions [**10**]. Specifically, the in-degree distribution exhibits a significantly heavier tail than the out-degree distribution. This proves that user visibility concentrates within a small elite, directly validating the rich-get-richer topology operating on the platform [**10**].

The infrastructure of the platform plays a decisive role in shaping visibility. The algorithms used for content curation, such as the "What's Hot" feed, act as multipliers for established popularity.

This creates a "Rich-Get-Richer" effect, where visibility generates more visibility. The data on interaction distribution is stark [15]:

- The network topology is characterized by a high concentration of engagement. A small minority of "hubs" attracts the vast majority of interactions, leaving most users in the periphery.
- This concentration proves that the network topology privileges a few central "hubs", replicating the power imbalances seen in centralized social networks.

2.4.4 The Algorithmic Glass Ceiling

A crucial theoretical finding in this analysis is the presence of an "Algorithmic Glass Ceiling". This phenomenon explains why certain groups struggle to gain visibility despite the platform's openness.

Recommendation systems, often based on "random walks" and "People you may know" features, are not neutral. They tend to reinforce existing biases through two main mechanisms:

1. **Amplification of Homophily:** Users tend to connect with people similar to themselves (homophily). Algorithms detect this pattern and suggest more of the same connections.

2. **Inverted Hierarchies:** It has been observed that male accounts often show much stronger homophily compared to female accounts. The algorithm picks up on this stronger cohesion and further promotes the majority group, effectively inverting representation hierarchies [15].

Mathematical models demonstrate that this algorithmic effect is systematically stronger than organic growth in generating disparity. It acts as an invisible barrier (glass ceiling) that exacerbates the underrepresentation of minorities at the top of the social hierarchy [15].

2.4.5 Conclusion of the Analysis

Bluesky represents a unique case study in the history of digital migrations. The platform provides unprecedented data transparency and allows users to control their algorithms. However, the analysis highlights a significant contradiction between the technical architecture and the social reality.

Despite its open design, Bluesky still suffers from the intrinsic problems of traditional social media. The data confirms two critical structural issues:

- **Concentration of Power:** The network topology is not distributed equally. Instead, it heavily favors a few dominant users. The analysis shows that the structure privileges a small number of "hubs" that attract the vast majority of interactions. This confirms that the "Rich-Get-Richer" dynamic replicates the hierarchies of centralized platforms, even in a decentralized protocol [2].

- **Ideological Echo Chambers:** The user base is politically homogeneous. While this significantly reduces the spread of fake news, it also creates "Asymmetric Polarization". Small minority groups are overwhelmed by the majority, creating an environment where diverse viewpoints struggle to emerge [5].

The theoretical implication is clear: technical freedom does not automatically guarantee social equality. Without specific corrections, the combination of human homophily (connecting with similar people) and algorithmic amplification tends to strengthen systemic biases. Therefore, the future challenge is to design algorithms that

include “rejection” or correction methods to prevent the structural marginalization of minority groups [15].

Chapter 3

Data and Methodology

3.1 Data Acquisition and Sampling Strategy

3.1.1 The AT Protocol Infrastructure: Firehose and Data Reachability

To understand the data acquisition strategy, it is necessary to analyze the federated architecture of the Authenticated Transfer Protocol (*atproto*). The network operates on a clear distinction between data hosting and data aggregation.

User data (posts, likes, feed definitions...) is stored in signed repositories hosted on independent Personal Data Servers (PDS). While these repositories ensure user sovereignty, they are fragmented across the network. To achieve global scale, the protocol utilizes Relays, which aggregate activity from all PDSs into a unified real-time stream known as the Firehose.

The Firehose is the primary primitive for data consumption in Bluesky. Accessible via WebSocket, it broadcasts authenticated events as they occur, allowing downstream services to sync with the network state. However, from a methodological perspective, the Firehose presents a limitation for static discovery. Being an event-driven stream, it only broadcasts updates and does not provide a queryable index of the historical state. To reconstruct the complete universe of existing custom feeds solely via the Firehose, one would need to observe the stream for an indefinite period, waiting for every single feed to generate a transaction.

3.1.2 Feed Ecosystem Extraction

Since the protocol architecture described above does not offer a centralized "List All" endpoint for custom feeds, we implemented a hybrid two-stage acquisition pipeline to overcome these discoverability limitations:

- **Phase 1: Discovery (Scraping).** We utilized 'Bluesky Directory' as a third-party aggregator to identify existing resources. A Python script was developed to iterate through the directory, filtering for specific URL path components (`bsky.app/profile/` and `/feed/`) to isolate valid feed identifiers. This process

allowed us to generate a static list of target URIs without relying on real-time event monitoring.

- **Phase 2: Retrieval (API).** The identified URLs served as input for the official data retrieval. Using the *atproto* SDK, we queried the App Views to access the authoritative metadata stored in the repositories. For each feed, the pipeline resolved the creator's unique Decentralized Identifier (DID) and fetched the core metrics (display name, description, creation timestamp, and total likes). Finally, an on-the-fly language detection was performed using the `langdetect` library.

3.1.3 Popularity Data Sampling

To analyze the dynamics of popularity and user behavior, we moved beyond the static definition of feeds to capture the active flow of interactions. Since scanning the entire user base (over 30 million registered accounts) was computationally prohibitive, we implemented a multi-stage sampling strategy designed to select a representative cohort of active users without introducing temporal or geographic biases.

The pipeline was executed in four logical steps

1. **Temporal Stratification and Real-Time Identification:** To capture a representative snapshot of active users while minimizing geographic bias, we implemented a scheduled monitoring system governed by a custom Python orchestrator. This controller managed the lifecycle of data acquisition sessions, launching independent subprocesses at four equidistant intervals throughout the day: 00:30, 06:30, 12:30, and 18:30 UTC.

   To interface with the network, the listener script instantiated a `Firehose SubscribeReposClient` from the `atproto` library, opening a direct channel to the global event stream. The listener script was configured to subscribe to the `com.atproto.sync.subscribeRepos` WebSocket endpoint for a fixed duration of 30 minutes per session.

   Crucially, the system performed session-scoped deduplication. Upon startup, each worker initialized an empty in-memory hash set. As new commit events arrived, the system extracted the repository DID and checked it against this local registry. Only DIDs not yet seen within the current 30-minute window were recorded. Following the completion of the daily cycle, an aggregation script concatenated the session logs into a unified corpus. From this aggregate, duplicates across different time windows were resolved, resulting in a final raw pool of approximately 500,000 unique active users.

2. **Short-Term Activity Crawl and User Profiling:** We harvested the full posting history of the identified pool for the 30 days prior to January 31, 2026. This phase was executed by a multi-threaded crawler (`crawltimelines_limited.py`) utilizing Python's `ThreadPoolExecutor` to process users in parallel batches.

To balance data completeness with collection speed, we established an upper bound on the data collection per user. Specifically, the crawler was programmed to halt the retrieval process once a user reached a total of 720 posts in the 30-day period. This approach capped the dataset size based on aggregate volume rather than daily frequency, allowing for natural fluctuations in user activity .

To circumvent the strict API rate limits (3,000 requests/5 min), the workload was distributed across four distinct worker instances, each managing a specific partition of users and operating under a unique authenticated session.

In a complementary phase, we executed a bulk profile enrichment script. Sociometric attributes (follower/following counts, creation date and bio) were retrieved via the `app.bsky.actor.get_profiles` endpoint. Leveraging the vectorization capabilities of the AT Protocol, these lookups were optimized into batched requests of 25 actors per API call, significantly reducing network overhead compared to sequential querying.

3. **Log-Transformed Stratified Sampling:** To select a representative sample of 30,000 users, we applied a data-driven stratification:

   - *Eligibility Filtering:* Inclusion required account age > 180 days, minimum 5 posts, and non-zero engagement.
   - *Logarithmic Transformation:* We calculated the Average Engagement per Post for all eligible users defined as:

   $$E_{avg} = \frac{Likes + Reposts + Replies}{PostCount} \tag{3.1}$$

   Recognizing that social media engagement follows a heavy-tailed Power Law distribution, we applied a logarithmic transformation:

   $$X = \log(1 + E_{avg}) \tag{3.2}$$

   This operation was necessary to normalize the feature space before clustering, ensuring that extreme outliers (typical of viral content) did not disproportionately bias the subsequent classification.
   - *Clustering:* We used K-Means ($k = 3$) to identify Low, Medium, and High engagement tiers, performing proportional sampling to maintain the ecosystem's natural proportions.

4. **Longitudinal Data Retrieval (The "Deep" Dataset):** For the final cohort of 30,000 users, we executed a deep historical crawl to reconstruct the long-term evolution of popularity dynamics. Extending the multi-threaded architecture established in the short-term analysis, this phase involved retrieving the complete timeline of activities covering the 180-day observation window.

The extractor iterated endpoint using a cursor-based pagination loop, stopping only when the event timestamp predated the six-month horizon.

To ensure computational feasibility, a cumulative cap of 4,320 items per user was enforced as a hard-stop condition within the retrieval loop. This threshold, designed to maintain the average frequency limit of approximately 24 daily actions, effectively bounded the execution time per user.

A critical component of this pipeline was the parsing logic developed to distinguish between content creation and propagation. By inspecting the `reason` attribute of the returned `FeedViewPost` objects, the system programmatically differentiated between the *Actor* (the user performing the action, extracted from `reason.by.did` in case of reposts) and the *Original Author* (the content creator). This distinction is fundamental to modeling the directed network of influence. The collected data was serialized in real-time into compressed JSONL files (`.jsonl.gz`) to minimize I/O latency, concluding with a parallel retrieval of updated sociometric attributes to capture the users' final reputation status.

3.2 Dataset Characterization

3.2.1 The Feed Landscape Dataset

This dataset constitutes a comprehensive census of the custom algorithmic ecosystem within the AT Protocol. It captures a static snapshot of 79,059 distinct feed generators, developed and published by a community of 33,901 unique creators over a period extending from May 2023 to January 15, 2026.

The collection aggregates metadata from multiple indexing sources, providing a structural overview of how algorithms are described, categorized, and adopted by the user base. The dataset organizes the attributes of each feed into three logical categories: Metadata (content description), Metrics (popularity indicators), and Identification (technical pointers). Table 3.1 details the schema used for the analysis.

3.2.2 The User Interaction Dataset

The "Deep" dataset represents the core resource for the longitudinal analysis of popularity dynamics presented in Chapters 5. We constructed the final dataset by performing a deep historical crawl on the stratified sample defined in Section 3.1.3. Of the initial target of 30,000 users, the pipeline successfully retrieved complete timelines for 29,993 unique profiles. For each of these users, we downloaded the full history of activities covering a 180-day observation window, ending on February 3rd. The resulting collection contains a total of 31,248,323 interactions, including both original posts and reposts.

To correctly model influence dynamics, the dataset strictly distinguishes between content creation and propagation. Each interaction identifies two key entities: the

Table 3.1: Feed Landscape Variables

Variable Name	Data Type	Description
name	String	The display name of the feed (e.g., "Science", "Black Sky").
description	String	A textual explanation of the algorithm's curation logic.
language	String	The detected dominant language of the feed's metadata.
creation_date	Datetime	The timestamp of the feed's initial publication.
feed_likes	Integer	The total number of likes received by the feed generator itself.
creator_did	String	The unique Decentralized Identifier of the developer.
creator_followers	Integer	The follower count of the feed's creator, used to correlate developer reputation with feed success.
uri	String	The precise AT Protocol resource identifier (at://...).

Actor (who performs the action) and the *Original Author* (who created the content). In standard posts these roles coincide, whereas in reposts the Actor redistributes the Author's content, establishing a directed connection in the social graph. This distinction is necessary to reconstruct the directed network of influence. The collected variables are detailed in Table 3.2.

Table 3.2: User Interaction Variables

Variable Name	Data Type	Description
did	String	Unique ID of the user performing the action.
original_author_did	String	Unique ID of the content's original creator.
created_at	Datetime	Exact timestamp of the action.
is_repost	Boolean	Flag indicating the nature of the action (True = Propagation, False = Creation).
uri	String	Unique resource identifier of the content.
text	String	The textual content of the post.
like_count	Integer	Number of likes received by the content at the time of extraction.
repost_count	Integer	Number of times the content has been shared.
reply_count	Integer	Number of direct replies to the content.
followers_count	Integer	The actor's follower count at the time of extraction.

To prepare the data for the mathematical modeling, we converted the raw collection into a linear chronological sequence. This step was essential to simulate the flow of time required by the stochastic equations, reorganizing millions of independent interactions into a single, ordered timeline.

This transformation involved three specific operations:

- **Chronological Ordering:** We aligned every interaction by its exact timestamp (ts) to reconstruct the linear progression of the ecosystem.
- **Metric Calculation:** We computed the total engagement v_i(defined as the sum of likes, reposts, and replies) explicitly distinguishing between original

posts and reposts.

- **Data Streamlining:** We reduced the complex data structure to essential numerical vectors (Timestamp, User ID, v_i, Followers), discarding metadata to maximize computational efficiency.

The final output of this transformation is a streamlined numerical matrix designed for high-performance stochastic simulation. Unlike the raw JSON logs, which contain textual metadata and nested structures, this processed dataset focuses exclusively on the temporal and quantitative dimensions of popularity. The resulting schema, detailed in Table 3.3, constitutes the empirical input for the Stochastic Differential Equations (SDEs) discussed in Chapters 5.

Table 3.3: Mathematical Model Input Schema

Variable	Type	Description
`ts`	Float	Continuous time variable (Unix Timestamp), representing the exact moment of the interaction t.
`did`	String	Unique identifier of the user performing the action.
`v_i`	Integer	Aggregate engagement volume, calculated as the sum of popularity metrics ($Likes + Reposts + Replies$).
`post_type`	Categorical	Classification of the event derived from the `is_repost` flag (0 = Original Creation, 1 = Propagation).
`followers_count`	Integer	The user's network size at the time of extraction, serving as a proxy for reputation potential.
`tier`	String	The user's activity class (Low, Medium, High) assigned via K-Means clustering.
`user_created_ts`	Datetime	The account creation timestamp, used to normalize user longevity.

This structured format allows us to treat the popularity dynamics as a continuous-time signal $X(t)$, enabling the direct calibration of the drift (γ) and volatility (ϕ) parameters from the event stream.

3.2.3 Handling Saturated Timelines and Data Truncation

As discussed in the data retrieval section, the collection pipeline included a hard limit of 4,320 interactions per user. This limit was necessary to keep the data processing manageable. However, it caused an issue for highly active users. When a user reaches the limit, the data collection stops. As a result, the oldest posts of these "saturated" profiles are missing from our dataset.

We could not keep these incomplete timelines in our final dataset. Including them would bias our mathematical models. In particular, truncated data negatively affects the calibration of the Stochastic Differential Equations (SDEs) for two reasons:

- **Distortion of the activity rate (λ):** Since the long-term history of these power users is restricted to an artificially shortened time window, their baseline posting frequency would be overestimated.

- **Missing initialization phase ($X(t)$):** The continuous-time decay function used to model popularity requires a clear point of origin to establish a steady-state baseline. For some saturated users, the initial "cold start" phase is unobserved, making it impossible to correctly initialize the model parameters.

To address this issue without completely discarding valuable network information, we decided to separate the saturated cohort ($N = 2,337$) from the main analysis.

These users were excluded from the primary dataset used for the dynamic stochastic simulations, ensuring that the remaining profiles have fully observable, untruncated timelines. However, since this hyper-active group represents a critical mix of genuine power users and automated bots, we retained their data for a dedicated exploratory analysis. The classification and behavioral study of this specific cohort is detailed in Chapter 4 (Section 4.2.3).

3.3 Data Processing Pipelines

3.3.1 Natural Language Processing (NLP) for Feeds

To analyze the semantic landscape of the custom algorithms, we implemented a dedicated NLP pipeline using BERTopic, a topic modeling technique that leverages transformers and c-TF-IDF to create dense clusters of interpretable text.

Preprocessing Protocol

Before modeling, the raw metadata underwent a strict filtering process to ensure quality and coherence:

- **Language Consistency:** We selected only feeds flagged as English (`lang='en'`) to align with the pre-trained capabilities of the embedding model.
- **Content Availability:** Entries with empty or null descriptions were removed to prevent noise in the vector space.
- **Temporal Windowing:** We excluded feeds created before December 31, 2022, filtering out legacy artifacts from the protocol's beta testing phase to focus on the current ecosystem.

Embedding and Clustering

The core modeling pipeline utilized the `all-MiniLM-L6-v2` sentence transformer to convert feed descriptions into dense numerical vectors. To manage the high dimensionality of the embeddings, we applied UMAP (Uniform Manifold Approximation and Projection) with a fixed random state (`random_state=42`) for reproducibility.

Subsequently, HDBSCAN was employed for density-based clustering, allowing the model to identify topics of varying shapes and sizes without forcing a predetermined number of clusters. Finally, the c-TF-IDF score was calculated to extract the

most representative keywords for each cluster, enabling the semantic labeling of the discovered topics.

3.3.2 Mathematical and Stochastic Modeling

To prepare the dataset for the popularity dynamics analysis presented in Chapters 5, we transformed the raw interaction logs into continuous temporal metrics. While the theoretical framework relies on Stochastic Differential Equations (defined in detail in Section 5.1.1), this preprocessing phase focused on operationalizing the key variables required by the model: Attention (X), Viral Capacity (θ), and Intrinsic Merit (β).

1. **Reconstruction of Attention Trajectories (X):** We modeled popularity as a dynamic stock of collective attention, converting discrete interaction timestamps into a continuous signal $X(t)$ via a numerical decay function. To mitigate the "cold start" bias caused by the arbitrary beginning of the observation window, we implemented a 30-day initialization phase. Data from this first month was used solely to prime the model, allowing the attention variable to reach a realistic steady-state level before commencing statistical estimation on the subsequent activity.

2. **Parameter Estimation via Grid Search:** Instead of estimating the preferential attachment exponent θ and the temporal decay γ through direct regression, we adopted a *Grid Search* optimization strategy. We explored the parameter space (γ, θ) by iteratively transforming the empirical engagement data according to the proposed differential equation model. For each combination of parameters, we evaluated the quality of the fit using the Kolmogorov-Smirnov (KS) distance and Log-Likelihood ratios. This allowed us to identify the specific (γ, θ) pair that best explains the observed popularity dynamics by minimizing the error between the theoretical distribution and the transformed data.

3. **User Profiling: Merit (β) and Activity (λ):** Finally, we derived the specific behavioral parameters for each of the 29,993 users:

    - **Intrinsic Merit (β_i):** Calculated as the user's ability to generate engagement normalized by their current popularity. Crucially, this metric was computed strictly on original posts (excluding reposts) to capture creative quality rather than curation habits.
    - **Activity Rate (λ_i):** We measured the temporal distribution of posting events, calculating the Dispersion Index (D) to verify the "bursty" nature of human activity ($D \gg 1$) compared to random Poisson processes.

The resulting parameter set constitutes the empirical input for the stochastic simulations discussed in Chapters 5.

Chapter 4

Bluesky - Exploratory Data Analysis

4.1 The Feed Ecosystem Analysis

4.1.1 Dataset Definition and Evolution

The analysis focuses on the population of $N = 79,059$ distinct feeds identified via the retrieval pipeline described in Section 3.1.2. This sample represents the cumulative output of the algorithmic marketplace from its inception until January 15, 2026.

To understand the dynamics of this ecosystem, we examine the temporal distribution of development efforts. As illustrated in Figure 4.1, the creation of new feeds does not follow a linear accumulation pattern but exhibits a highly non-stationary behavior characterized by distinct "bursts" of activity.

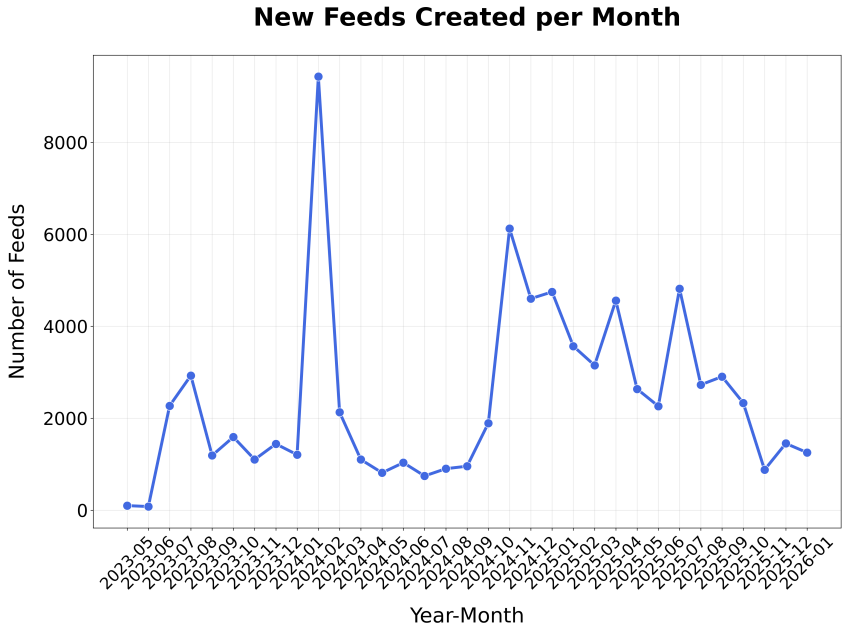

Figure 4.1: Temporal distribution of feed creation and development bursts.

The time series reveals that the supply of algorithms is reactive to external

adoption shocks rather than organic internal growth. Two major structural breaks are observable:

- **February 2024:** A sharp spike reaching over 9,000 new feeds per month. This correlates with the removal of the invite-only system, which opened the platform to the general public and triggered an immediate demand for content curation tools.
- **October 2024 (Q4):** A secondary resurgence in development activity (approximately 6,000 feeds), likely driven by the second wave of user migration from competing platforms (X/Twitter). This influx was primarily catalyzed by the platform's localized ban in Brazil and subsequently reinforced by the US electoral period.

Between these events, the ecosystem shows a baseline production rate that stabilizes at lower levels, suggesting that developer enthusiasm is strongly coupled with sudden influxes of new users.

Regarding the linguistic and geographic composition, the metadata analysis confirms a significant hegemony of the English language, which characterizes approximately 57% of the identified feeds. However, the ecosystem is not monolingual. Outside the English-dominant core, a substantial development cluster has emerged in Japanese, representing 22% of the ecosystem.

4.1.2 Creator Dynamics and Concentration of Supply

The supply side of the algorithmic marketplace is driven by a population of $N = 33,869$ unique creators. However, the contribution of these actors is highly asymmetrical. By analyzing individual productivity (defined as the number of distinct feed generators published by a single user) we observe a distribution that deviates strongly from a normal curve, following a heavy-tailed power law.

As illustrated in Figure 4.2, the ecosystem is dominated by a "long tail" of casual developers. The data indicates that 62.3% of the population (21,089 creators) has published exactly one feed. This suggests that for the majority of users, the barrier to entry is low enough to experiment with the feature, but the incentive to maintain a portfolio of algorithms is lacking. Conversely, the right tail of the distribution decays rapidly, indicating that high-volume production is reserved for a specific subset of actors.

To understand the nature of these high-volume producers, we compared the top creators ranked by output volume against those ranked by popularity. This comparison reveals a stark "Quantity vs. Quality" paradox within the ecosystem.

Figure 4.3 displays the Top 20 creators by the number of feeds published. The ranking is led by entities such as `@bluestocks.app` and `@trending.bsky.app`, which have generated over 2,500 and 2,000 feeds respectively. A closer inspection of the engagement metrics for these accounts reveals the limits of this strategy: despite their massive output, the average engagement per feed is often negligible. This pattern

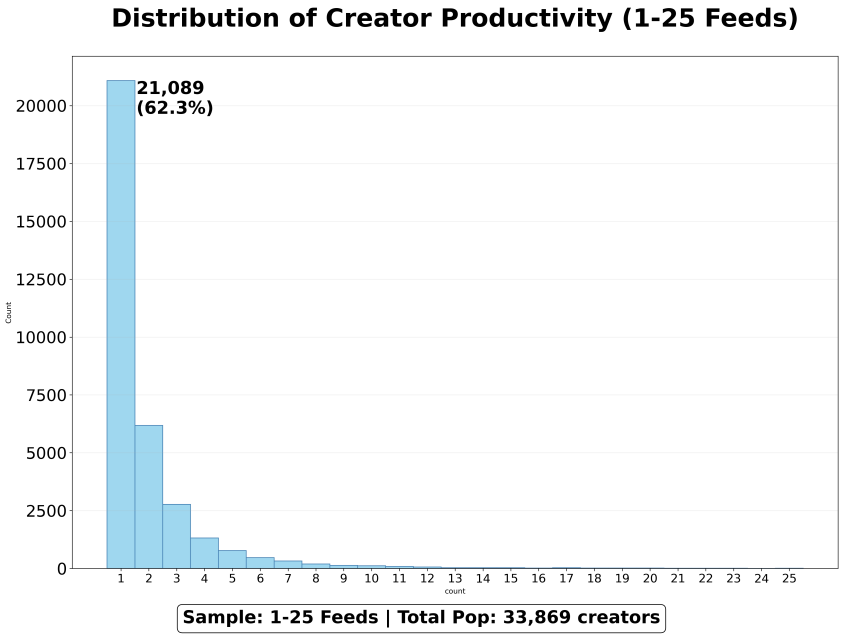

Figure 4.2: Distribution of creator productivity (Power Law).

points to the presence of automated systems or "bot-like" behaviors designed to flood the marketplace with programmatic variations of content filters.

Figure 4.4, conversely, ranks the Top 20 creators by cumulative likes. Here, the hierarchy is defined by institutional authority and tool utility rather than volume. Notably, there is zero overlap between the top 5 producers of feeds and the top 5 most popular creators.

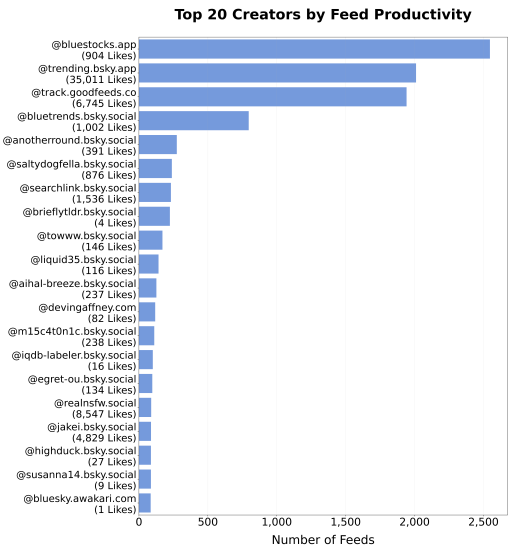

Figure 4.3: Top 20 creators by the number of feeds published.

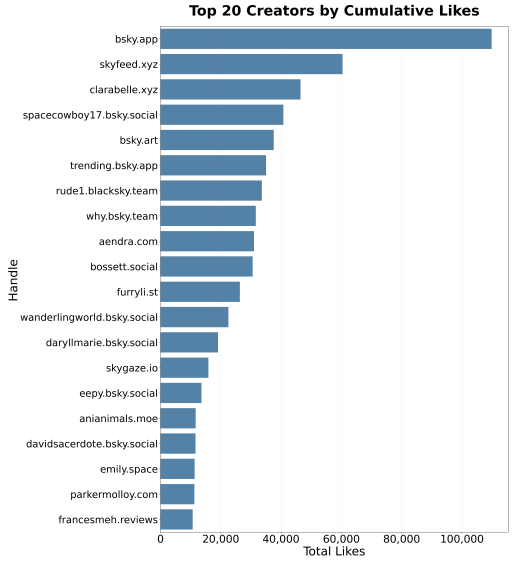

Figure 4.4: Top 20 creators by cumulative likes.

4.1.3 Popularity and the Attention Gap

While the supply of algorithms is abundant, the distribution of user attention is characterized by extreme scarcity. By analyzing the aggregate volume of likes received by each of the 79,059 feeds, we can identify the primary filter in this ecosystem: the "Zero-Like Phenomenon."

As shown in Figure 4.5, a substantial portion of the developed algorithms fails to overcome the initial discovery barrier. Specifically, 45.8% of the total population (36,160 feeds) has never received a single like. If we include the stratum of feeds that have received exactly one like, the failure rate rises to 63%. This indicates that the majority of the marketplace exists in a state of dormancy, technically available but functionally invisible.

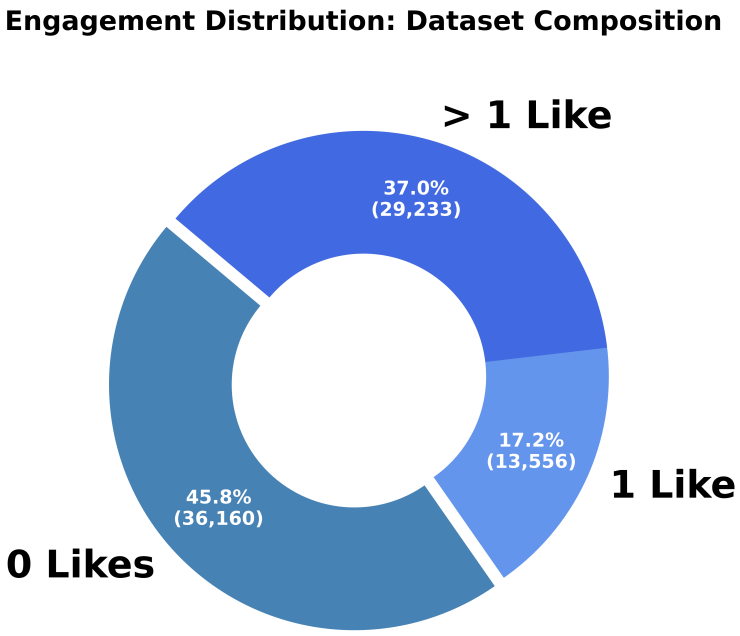

Figure 4.5: The Zero-Like Phenomenon and discovery barriers.

For the surviving 37% of feeds that manage to accrue engagement, the distribution of popularity follows a steep stratification. Figure 4.6 categorizes the feeds into logarithmic buckets of success. The data reveals that 90.2% of all feeds remain trapped in the lowest tier of popularity (0-10 likes). The probability of a feed migrating to a higher tier decreases exponentially: only 8.1% reach the moderate tier (11-100 likes), while the "viral" threshold (>1,000 likes) is breached by a mere 0.2% of the population.

To understand the characteristics of the "Super-Elite," we analyzed the Top 20 most popular feeds in Figure 4.7. The ranking confirms the persistence of institutional hegemony: the list is led by "Popular With Friends," an algorithm developed by the official `bsky.app` account. However, the presence of third-party algorithms such as "Mutuals" and specific topical feeds like "Science" demonstrates that non-institutional actors can compete if they provide significant utility.

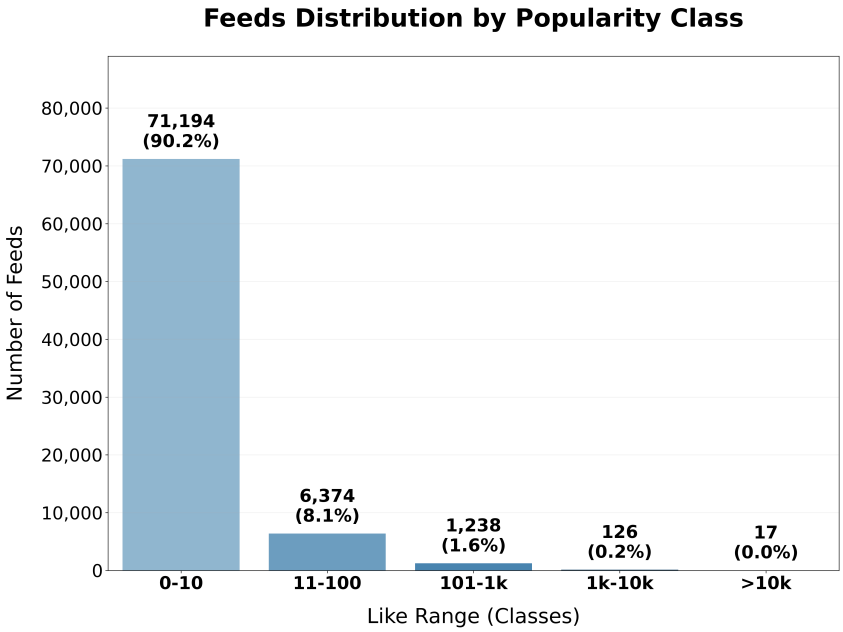

Figure 4.6: Logarithmic stratification of feed popularity.

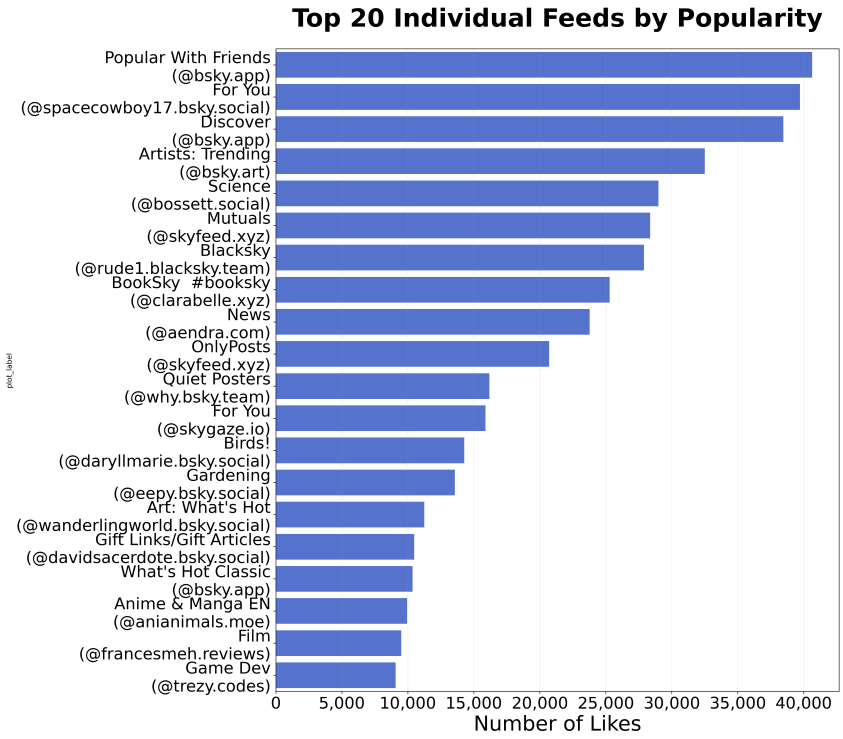

Figure 4.7: Top 20 most popular feeds by total likes.

4.1.4 Determinants of Success: Longevity and Social Capital

Given the extreme stratification of popularity, it is crucial to determine whether success is driven by structural advantages such as the creator's pre-existing reputation (Social Capital) or the time of entry into the market (Longevity).

The Role of Social Capital

The data reveals a clear disconnect between a creator's social popularity (follower count) and the actual success of their custom feeds (feed likes). As illustrated in Figure 4.8, the correlation between a creator's follower count and their feed's popularity is negligible ($r = 0.08$) for 99.5% of the population.

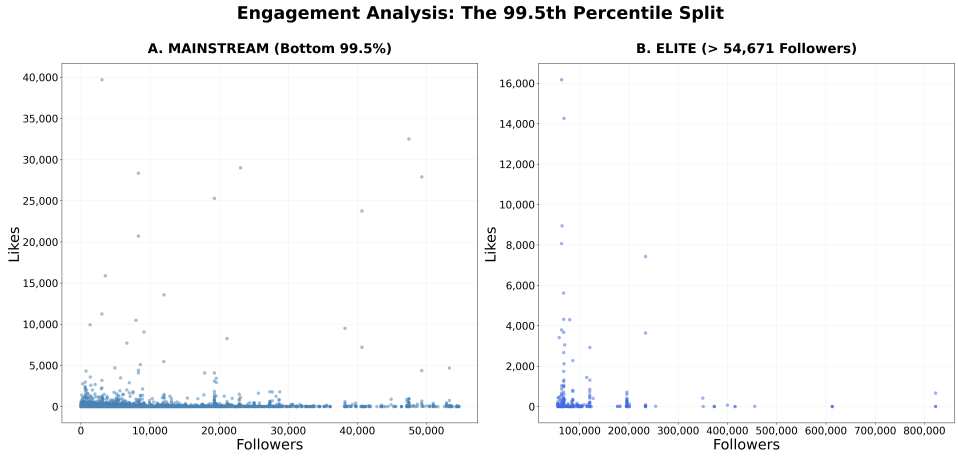

Figure 4.8: Correlation analysis between Social Capital and feed popularity.

However, while social capital is not a sufficient condition for success, it acts as a structural constraint. The stratification analysis in Figure 4.9 demonstrates a "glass ceiling" effect: while "Star" creators (>10k followers) can reach peaks of >30k likes, "Emerging" and "Micro" creators rarely exceed the 1,000-like threshold.

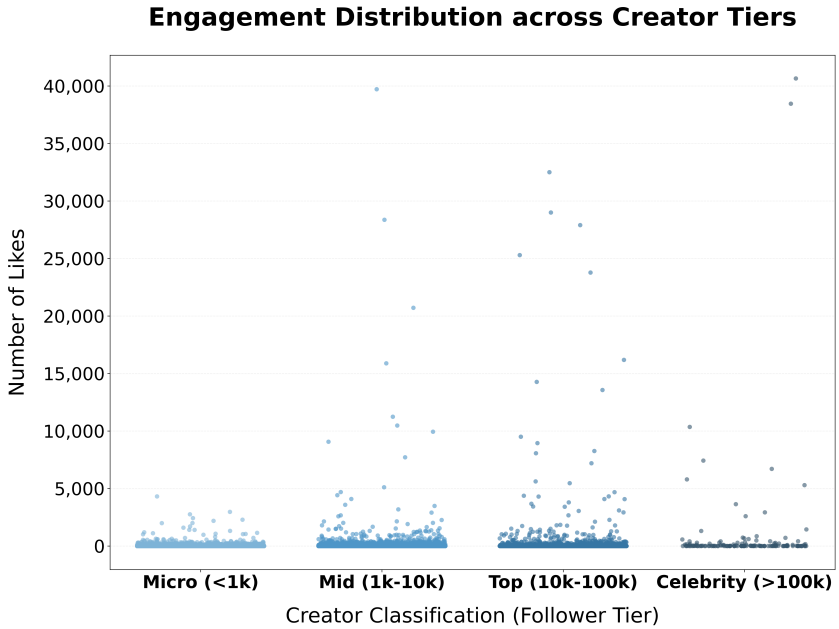

Figure 4.9: The glass ceiling effect in algorithmic success.

The Role of Longevity

The analysis refutes the existence of a "First-Mover Advantage." Figure 4.10 plots the popularity of feeds against their age, showing a flat regression line that indicates zero correlation between longevity and success. High-performing feeds are distributed uniformly across the temporal axis, confirming that the ecosystem remains dynamic and allows new entrants to displace incumbents based on utility rather than seniority.

This pattern highlights a fundamental difference between social capital and algorithmic popularity. While human follower counts typically exhibit monotonic growth over time (driven by the cumulative nature of social networks [3]) feed likes are governed by a logic of continuous utility. Users appear to treat feed subscriptions as active choices: if an algorithm loses its relevance or a user's interests shift, the like is often revoked. This "unliking" mechanism prevents legacy feeds from accumulating popularity through mere temporal inertia, ensuring that the marketplace remains contestable and that success is tied to current performance rather than historical seniority.

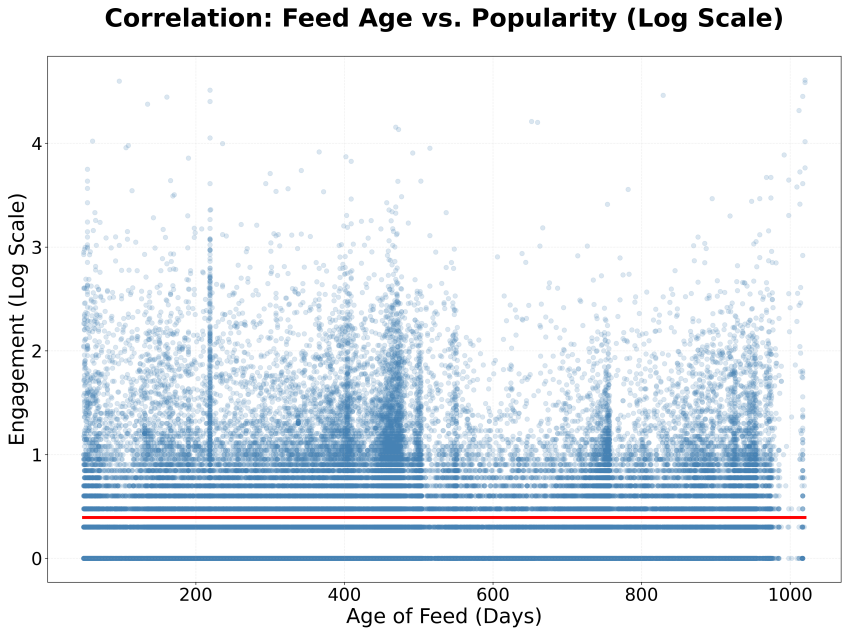

Figure 4.10: Longevity vs. Popularity.

4.2 The User Interaction Dataset

The analysis presented in the subsequent chapters relies on a stratified sample of $N = 29,993$ unique users, selected to represent the active core of the Bluesky ecosystem.

To ensure the statistical robustness of the behavioral models, the dataset underwent a strict filtering process. Specifically, we excluded inactive accounts by enforcing a minimum requirement of 5 posts within the 180-day observation window (August 2025 – January 2026). This threshold ensures that the analysis focuses solely on

users who actively drive the interaction dynamics, effectively filtering out passive observers and dormant profiles. Despite this selectivity, the resulting dataset remains extensive, capturing a total volume of over 31 million interactions.

From a demographic perspective, the sample reflects the current linguistic footprint of the network. As expected, the ecosystem is predominantly Anglophone, with English accounting for the vast majority of the content volume [21]. Japanese emerges as the second-largest linguistic block, standing out as a significant and self-contained community within the dataset. Following this, smaller but active clusters of German, Spanish, and Portuguese speakers are observable. This linguistic distribution suggests that while the platform is technically global, the interaction dynamics modeled in this thesis are primarily driven by Western and Japanese behavioral patterns.

4.2.1 Stratification Strategy and Behavioral Clustering

To ensure statistical representativeness, we employed a proportional stratified sampling method based on user engagement. Starting from an active candidate population (filtered for minimum activity of 5 posts in 180 days), we applied a K-Means clustering algorithm ($k = 3$) on the log-transformed average engagement to identify natural behavioral classes.

The resulting segmentation is illustrated in Figure 4.11. The algorithm established specific cut-off thresholds that strictly define the three tiers:

- **Tier 1 (Low Engagement):** Users with $\bar{E} < 3.4$. This group constitutes the majority (66.8%), representing the standard user base.
- **Tier 2 (Medium Engagement):** Users with $3.4 \leq \bar{E} < 33.6$. A transitional class (27.1%) of consistently active participants.
- **Tier 3 (High Engagement):** Users with $\bar{E} \geq 33.6$. The "Elite" tail (6.1%), characterized by viral interactions.

The final sample of $N = 29,998$ users was extracted to preserve these exact proportions, ensuring that the statistical properties of the sample mirror those of the full population (Table 4.1).

Table 4.1: Population vs. Sample Stratification Statistics

Tier	Pop. Count	Sample Count	Share (%)
Low	54,775	20,033	66.8%
Medium	22,272	8,145	27.1%
High	4,977	1,820	6.1%
Total	**82,024**	**29,998**	**100%**

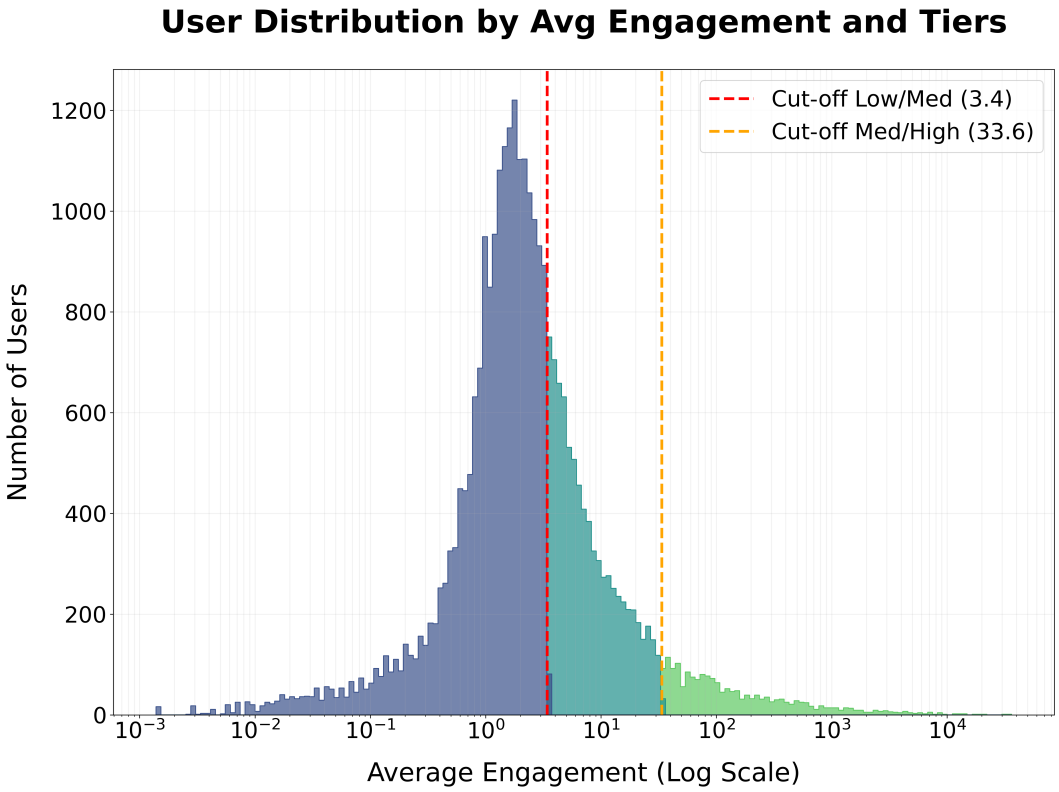

Figure 4.11: Distribution of users by Average Engagement. The vertical dashed lines indicate the K-Means cut-offs separating the three tiers.

4.2.2 Activity Distribution and Data Saturation

The distribution of user activity, defined as the combined volume of posts and reposts, is shaped by two main factors. First, it reflects a social stratification typical of complex networks, where a small minority of users generates the vast majority of content. Second, the data is affected by a technical truncation (or "temporal censorship") caused by the retrieval architecture; since the API limits the history to a fixed number of posts, the older activities of the most prolific users are effectively cut off.

The Power Law of Participation

First, we analyze the aggregate volume of activity per user. As illustrated in Figure 4.12, the population follows a heavy-tailed "Power Law" distribution. The median user produces content sporadically (typically fewer than 50 posts in six months) while a hyper-active minority drives the vast majority of the conversation volume.

Notably, the distribution exhibits a distinct artifact at the extreme right boundary. The bin adjacent to the retrieval limit ($N = 4,320$) shows a disproportionate spike in frequency. This accumulation is not a behavioral feature but a structural consequence of the API constraint: the natural tail of the distribution would extend asymptotically beyond this point, but all users exceeding the cap are compressed into this single terminal category.

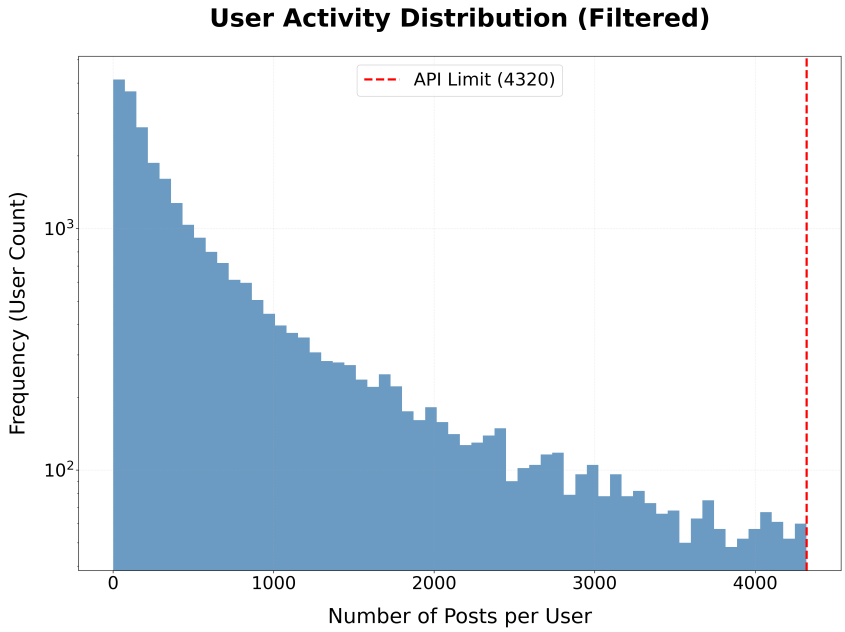

Figure 4.12: Distribution of user activity.

Temporal Saturation and Bias

Consequently, for the hyper-active users identified in the tail, the retrieval limit imposes a strict boundary on historical visibility. This technical constraint introduces a temporal bias that is visually quantified in Figure 4.13.

The graph compares the daily volume of posts (grey bars) against the availability of "Power Users" in the dataset (red dashed line). As we observe the timeline backwards, the red curve declines significantly faster than the total volume. This divergence confirms that while standard users (92% of the sample) are tracked for the full 180 days, the history of high-frequency users is progressively lost due to buffer saturation.

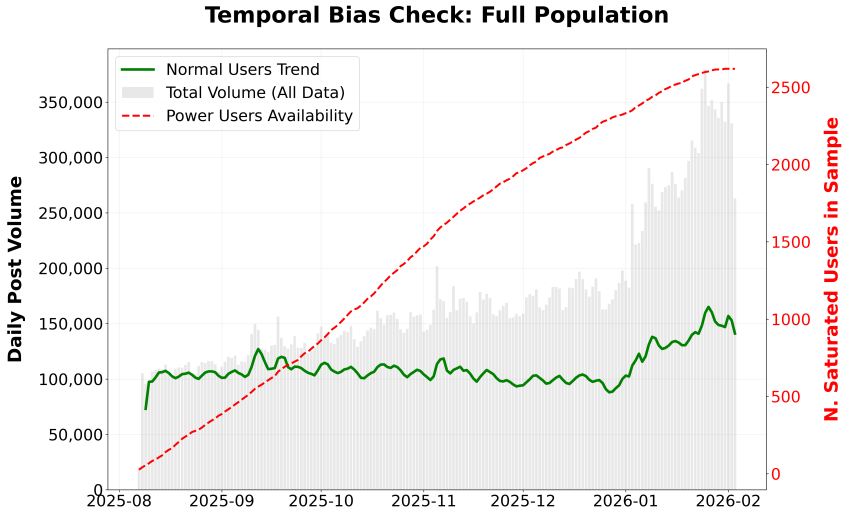

Figure 4.13: Analysis of temporal bias.

4.2.3 Anatomy of the Saturated Class: Bot Filtering

The "Saturated Class" ($N = 2,337$), comprising users who hit the API retrieval limit, represents a critical intersection of genuine "Power Users" and automated agents. To prevent data pollution, we implemented a multi-dimensional heuristic classifier designed to disentangle human actors from algorithmic scripts.

The classification relied on four behavioral metrics: *Temporal Entropy* (to detect rigid scheduling), *Originality Ratio*, *Content Duplication*, and *Engagement Quality*. By mapping users into this feature space, we applied a strict decision tree:

- **Bot Identification:** Users were flagged as automated if they exhibited high content duplication ($D_{max} > 3$), negligible originality ($R_{orig} < 5\%$), or consistently low impact ($\bar{L} < 2.0$).
- **Human Verification:** Users demonstrating high temporal irregularity ($H_t > 5.0$) were confirmed as active humans.

As illustrated in Figure 4.14, this approach revealed a clear separation in the dataset. The analysis classified **60.6%** of the saturated population ($N = 1,417$) as **Bots**, predominantly belonging to the Low Engagement Tier (spam/noise). Conversely, **38.6%** ($N = 902$) were identified as legitimate **Humans**, while a negligible fraction (0.8%) remained ambiguous.

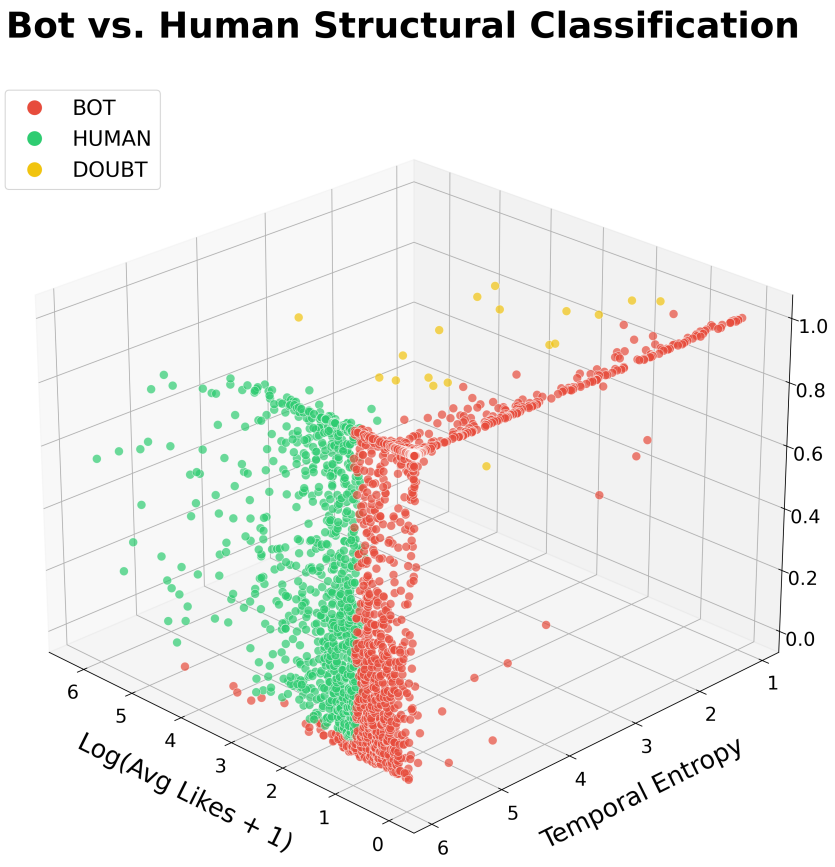

Figure 4.14: 3D Classification of the Saturated Class.

To validate these results, we performed a manual inspection on a stratified random sample of 20 profiles per category. The qualitative review confirmed the robustness of the classifier:

- The **Human** cluster showed zero false positives (no automated behavior detected).
- The **Bot** cluster was confirmed to consist primarily of automated scripts. Only 2 appeared to be human-operated; however, these users exhibited aggressive spamming behavior (high frequency, repetitive content) with negligible community engagement, effectively rendering them indistinguishable from algorithmic noise.
- The **Ambiguous** profiles ($N = 18$) were identified primarily as semi-automated news aggregators.

As explained in Chapter 3 (Section 3.2.3), we separated this saturated class ($N = 2,337$) from the main dataset. The analysis in this section confirmed that many of these hyper-active accounts are indeed automated bots.

However, even the genuine human users in this group have truncated data because of the API limit. Their oldest posts are simply missing.

Because of this, we decided not to consider this entire group for the next phases of our research. We did not delete their records, but we simply left these 2,337 users out of the longitudinal analysis. As a result, the following section on the "Cost of Attention" (Section 4.2.4) and the mathematical models in the next chapters focus exclusively on the remaining users. This ensures our final results are not biased by incomplete timelines or algorithmic spam. Specifically, the exclusion of these 2,337 profiles, each constrained by the 4,320, post limit yields a rigorously filtered dataset. The final cohort retained for the subsequent longitudinal analysis consists of 1,434 High-Engagement and 7,282 Medium-Engagement unique users, providing a reliable and untruncated baseline over a continuous observation window of 150 days.

4.2.4 The Cost of Attention: Activity vs. Popularity

A fundamental question for the popularity model is whether increased output leads directly to increased recognition. Figure 4.15 plots the relationship between user activity (Post count) and total engagement (Likes) on a log-log scale.

The data shows a clear positive correlation, with a regression slope of 1.26. This indicates that, on average, a higher volume of activity is associated with greater popularity. However, the scatter plot also reveals an high variance: for any given level of activity (e.g., 10^3 posts), the engagement received can vary by several orders of magnitude, ranging from a few dozen to millions of likes.

This dispersion confirms that while activity is a necessary driver, it does not guarantee success. The fact that users with identical post volumes achieve such different results suggests that popularity is governed by an intrinsic quality parameter

(defined in our model as "Merit" (β)) which acts as the true differentiator beyond mere quantity.

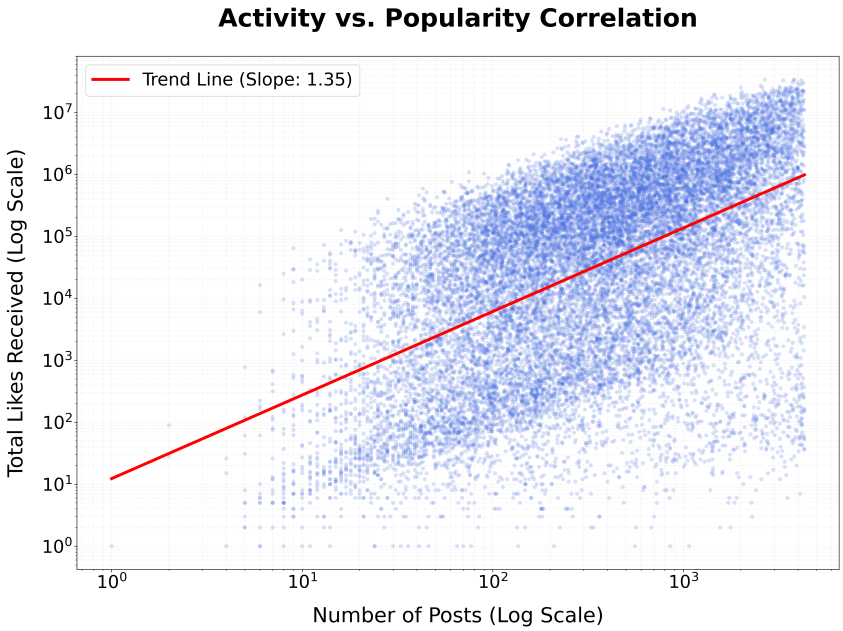

Figure 4.15: Relationship between total activity and average engagement.

Chapter 5

Global Dynamics of Popularity and Social Stratification

5.1 The Stochastic Framework: A Mean-Field Approach

Standard metrics such as follower counts are cumulative and rarely decrease, making them poor indicators of real-time influence. To capture the fluctuating nature of social media attention, we adopt a Mean-Field approach. Instead of modeling the complex web of pairwise interactions between users, this framework simplifies the system by analyzing how a single representative actor interacts with the aggregate "field" of collective attention. Accordingly, we define the state of an actor i at time t using the stochastic variable $X^{[i]}(t) \in \mathbb{R}^+$. We term this variable Effective Popularity: unlike static follower counts, $X^{[i]}(t)$ is a dynamic measure representing the instantaneous volume of attention the user commands in the ecosystem, subject to decay over time and sudden spikes driven by activity.

5.1.1 The Model Dynamics

The evolution of popularity is governed by a Stochastic Differential Equation (SDE) that captures the interplay between the natural decay of collective attention and the impulsive updates driven by content creation. Formally, the dynamics of $X^{[i]}(t)$ are defined by :

$$dX^{[i]}(t) = -\gamma X^{[i]}(t)dt + V_i(t)N_I^{[i]}(dt) + W_i(t)N_E^{[i]}(dt) \tag{5.1}$$

The equation comprises three distinct components:

1. **Physiological Decay ($-\gamma X^{[i]}(t)dt$):** In the absence of additional stimuli, popularity tends to decrease at a rate γ. This continuous decay models the limited nature of collective attention and represents the combined effect of users naturally losing or shifting interest, alongside the competition from other influencers. This formulation elegantly bypasses the need to explicitly model complex pairwise interactions among users; instead, γ captures the aggregate, mean-field effect of all competitors.

2. **Endogenous Jumps ($V_i(t)N_I^{[i]}(dt)$):** This term represents the change in popularity generated by the user's direct activity, namely the emission of a post. $N_I^{[i]}(dt)$ is a counting process triggered by new publications. Upon publication, the post attracts attention and creates a popularity "jump" modeled by the random variable $V_i(t)$. The magnitude of this jump cumulatively depends on mechanisms like algorithmic curation, user sharing behaviors, and exposure to the established follower base. Consequently, the expected success of a post depends on the user's current popularity $X^{[i]}(t^-)$ right before the publication, formalized as:

$$\mathbb{E}[V_i(t)|X^{[i]}(t^-) = x] = \epsilon + \beta_i x^\theta \tag{5.2}$$

   In this formulation, ϵ is a small positive constant introduced to avoid absorbing states. The parameter β_i is influencer-specific and reflects their intrinsic competence and ability to create captivating content. Conversely, the exponent θ captures the non-linear visibility amplification governed by user behavior and platform engagement algorithms, and is therefore assumed to be independent of the specific user. Furthermore, to capture the high variability of online success, the actual jump $V_i(t)$ is assumed to scale proportionally to its expected value, driven by a baseline random variable (empirically validated as a lognormal distribution). It is also worth noting that the frequency of these jumps can also be dependent on the user's effective popularity, reflecting how influencers adjust their posting behavior based on their current engagement.

3. **Exogenous Shocks ($W_i(t)N_E^{[i]}(dt)$):** Finally, this term accounts for external factors that can suddenly affect a user's popularity, such as newsworthy events reported by traditional media that expose the influencer to a broader audience. These external events cause popularity jumps denoted by $W_i(t)$. Because of their external nature, these jumps are considered reasonably independent of the user's behavior on the platform , and their arrivals, $N_E^{[i]}(dt)$, are modeled as a Poisson process. However, due to the lack of available information regarding these unpredictable exogenous events, this component is neglected in our empirical analysis, effectively setting $N_E^{[i]}(dt) = 0$.

5.1.2 Stationary Conditions and Ergodicity

To analyze the long-term distribution of popularity, it is necessary to determine under which conditions the system reaches a dynamic equilibrium. According to the theoretical analysis of the underlying Markov process, the popularity process $X(t)$ admits a unique stationary probability measure π only if the self-reinforcing mechanisms of the network are properly bounded. The model identifies two key exponents governing this feedback:

- θ: dictates that the expected popularity jump (likes) of a new post depends on the user's current popularity as x^θ.

- ϕ: implies that the user's posting rate intensity (behavior) depends on their current popularity as x^{ϕ}.

The primary sufficient condition for the system to be ergodic is defined by the inequality:

$$\theta + \phi < 1 \tag{5.3}$$

Interestingly, the theoretical framework demonstrates that the system can still reach a stable equilibrium even in the critical boundary case where $\theta + \phi = 1$. However, this is only possible provided that the generative parameters (such as the influencer's intrinsic ability and posting rates) are sufficiently small compared to the platform's physiological decay rate γ.

If these conditions hold, the system does not diverge. More importantly, the system's ergodicity implies convergence to this stationary distribution regardless of the user's initial conditions. This guarantees that the time averages of the popularity metrics will eventually converge to their expected values under the stationary distribution , whose exact shape can be mathematically described by an ordinary integro-differential equation.

From an applied perspective, this mathematical guarantee is crucial: it allows us to treat the empirical data collected in the 'Deep Dataset' not merely as a snapshot of a transient state, but as a robust and representative sample of the system's stationary regime.

5.1.3 Evolution of the Popularity Distribution

To describe the global distribution of popularity across the entire population, we move from the single-particle trajectory (SDE) to the macroscopic evolution of the probability distribution. According to Theorem 2 of the reference framework , the time evolution of the partition function $F(y,t) = \mathbb{P}(X(t) \leq y)$ is governed by the following Partial Integro-Differential Equation:

$$\frac{\partial F(y,t)}{\partial t} = \gamma y \frac{\partial F(y,t)}{\partial y} - \int [\lambda(t)\bar{F}_V(y-x|x) + \mu\bar{F}_W(y-x)] \frac{\partial F(x,t)}{\partial x} dx \tag{5.4}$$

In this formulation, $\bar{F}_V(\cdot|x)$ represents the conditional Complementary Cumulative Distribution Function (CCDF) of the endogenous popularity jumps V, while $\bar{F}_W(\cdot)$ is the CCDF of the exogenous shocks W. The parameter μ denotes the arrival rate of the exogenous events modeled by the Poisson process N_E.

Assuming the ergodicity conditions are met, in the stationary limit ($\partial F/\partial t = 0$), the distribution satisfies the corresponding Ordinary Integro-Differential Equation:

$$\gamma y \frac{dF(y)}{dy} = \int [\lambda(x)\bar{F}_V(y-x|x) + \mu\bar{F}_W(y-x)] \frac{dF(x)}{dx} dx \tag{5.5}$$

This equation bridges the gap between the microscopic stochastic behaviors of individual users (such as posting rate λ, intrinsic merit β, and the platform's decay γ)

and the macroscopic equilibrium of the network. Solving this equation is fundamental: it enables a probabilistic study of popularity emergence and allows us to predict the long-term stratification of the network, quantifying an influencer's probability of reaching and maintaining privileged positions.

5.1.4 Parameter Interpretation

The model relies on three key structural parameters that govern the network's dynamics and will be empirically estimated for the Bluesky ecosystem in Section 5.2:

- **Decay Rate (γ):** Represents the "metabolic cost" of attention. It quantifies how quickly the community forgets content, representing the rate at which popularity decreases in the absence of additional stimuli. Crucially, this decay is a consequence of the combined effect of the users' natural tendency to lose or shift interest, and the constant competition from other influencers on the platform.
- **Viral Capacity (θ):** This exponent ($0 \leq \theta \leq 1$) measures the strength of the "rich-get-richer" dynamic, determining the non-linear dependence of popularity jumps on current popularity. It is assumed to be primarily determined by users' collective behavior and the platform's engagement algorithms, making it independent of the specific influencer. A value of $\theta \approx 1$ amplifies the disparity among influencers, indicating a winner-takes-all market where current popularity heavily dictates future success. Conversely, a lower θ suggests a more egalitarian environment where visibility saturates.
- **Intrinsic Merit (β_i):** A user-specific scalar that determines the expected impact of a post, reflecting the influencer's overall ability. Unlike θ, which is a systemic property of the platform, β_i captures the specific ability of user i to generate engagement per unit of attention. It subsumes two distinct aspects: the intrinsic average quality of the content they produce, and the baseline level of engagement characteristic of their specific audience.

5.2 Empirical Calibration: The Bluesky Landscape

In this section, we apply the stochastic framework to the empirical data collected in the 'Deep Dataset'. Our objective is to estimate the specific parameters (viral capacity (θ) and decay rate (γ)) that drive the popularity dynamics within the AT Protocol.

5.2.1 The Optimization Landscape (Grid Search)

The core challenge in calibrating the SDE defined in Eq. (5.1) is that the system-level parameters (viral capacity θ and decay rate γ) are not directly observable in the raw data. To derive these values, we adopted the methodology proposed by Galante et al.

[3], implementing a Grid Search strategy based on Maximum Likelihood Estimation. The goal was to identify the parameter pair (γ^*, θ^*) that minimizes the statistical distance between the model's predictions and the observed user behavior.

The optimization algorithm operated iteratively over the parameter space. For each candidate pair (γ, θ), the procedure first reconstructed the continuous popularity history $X^{[i]}(t)$ for the analyzed user population. To ensure statistical significance, the initial pool refined to a robust cohort of 25,453 individuals through a strict preliminary filtering pipeline. Specifically, the algorithm excluded inactive users with fewer than 10 tracked posts, removed records with missing structural metadata, and discarded anomalous mathematical artifacts (such as NaN, infinite, or non-positive residuals) generated during the continuous-time integration. Finally, users for whom the Maximum Likelihood Estimation algorithm failed to converge were also excluded from the final dataset.

To comprehensively map the error landscape, the Grid Search explored a wide spectrum of configurations. The viral capacity θ was varied from 0.1 to 1.0 in increments of 0.1. The decay rate γ was explored by varying its reciprocal $1/\gamma$ (representing the system's inertia in days) on a logarithmic scale from 4 up to 2,048 days. To ensure accurate continuous-time integration between consecutive posts, the daily decay rates were mathematically converted into their corresponding per-second values ($\gamma_{sec} = 1/(\text{days} \times 86400)$) before being applied to the empirical timestamps. Furthermore, this entire parameter space was evaluated twice to account for different definitions of post success: once strictly measuring *likes*, and a second time assessing overall *engagement* (defined as the sum of likes, replies, and reposts). However, since preliminary tests showed that the results for overall engagement are almost identical to those obtained using only likes, the following sections will focus exclusively on the likes metric to avoid unnecessary repetition.

For the surviving population under each configuration, the procedure isolated the stochastic component of the dynamics by computing the normalized jump residuals $Z_{i,k} = V_{i,k}/(X^{[i]}(t_k))^{\theta}$, effectively decoupling the individual post performance from the user's accumulated status. The goodness-of-fit for each configuration was then quantified via the Kolmogorov-Smirnov (KS) distance.

The results visualized in Figure 5.1 reveals a distinct topology. The gradient of the cost function clearly converges towards a specific region of stability (represented by the red area), identifying a robust global minimum corresponding to high values of the viral capacity θ and a fast decay rate γ (i.e., low values of the characteristic time $1/\gamma$). This combination suggests a highly volatile ecosystem: while the strong rich-get-richer dynamics ($\theta \approx 1$) allow for massive, rapid concentrations of visibility, the extremely short collective memory ensures that attention fades quickly. Consequently, popularity is fleeting, and maintaining a dominant position requires continuous and frequent content generation to combat the platform's metabolic cost.

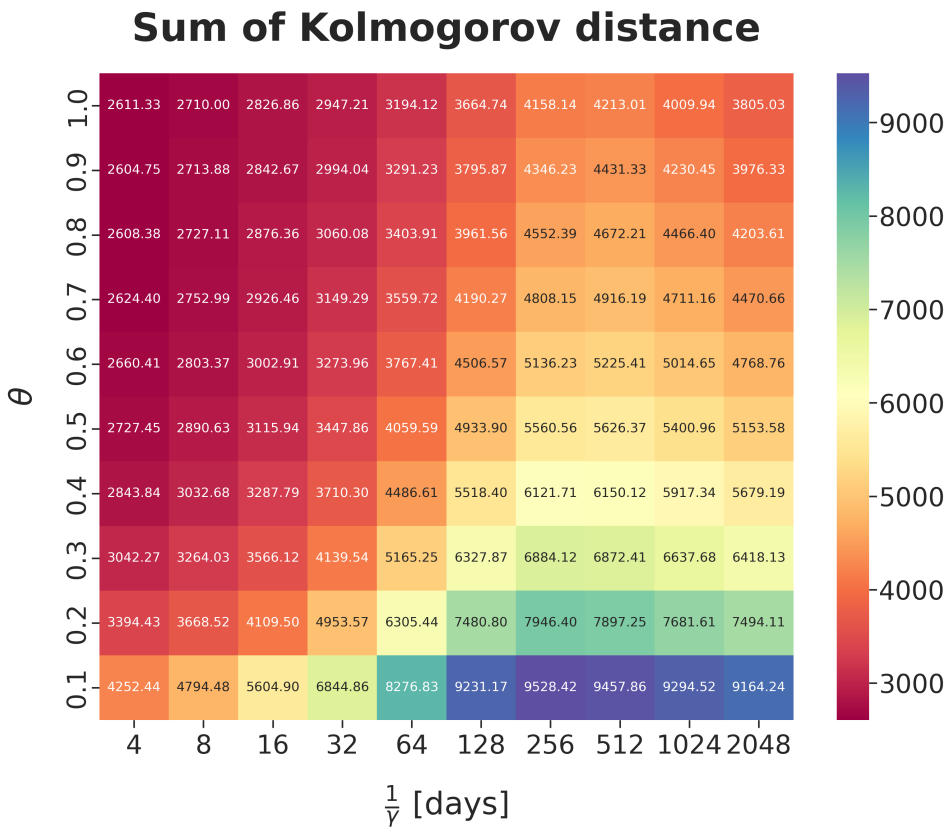

Figure 5.1: The Optimization Landscape (Heatmap). The color gradient represents the mean KS distance (lower values indicate better fit).

5.2.2 Global Parameter Estimation

The optimization process yielded a global minimum at ($\theta^* = 0.9, \tau^* \approx 4$ days). These parameters paint a picture of Bluesky as a high-velocity, high-virality environment, fundamentally different from some other platforms.

The Viral Exponent (θ) and System Stability

The optimal value for the viral capacity was empirically identified as:

$$\theta^* = 0.9 \tag{5.6}$$

A careful inspection of the optimization landscape reveals a plateau of high performance for values of θ approaching unity with low values of γ. In the heatmap displaying the *mean* KS distance 5.2, the configurations for $\theta = 0.7, 0.8, 0.9, 1.0$ appear numerically equivalent due to averaging artifacts. However, a deeper analysis of the *cumulative* error (the sum of KS distances across the entire population) in figure 5.1 resolves this ambiguity, showing that $\theta = 0.9$ yields a lower total discrepancy compared to the linear case.

Crucially, the "Count" heatmap (figure 5.3) confirms that the number of successfully fitted users is identical for both configurations. This consistency ensures that the performance advantage of $\theta = 0.9$ is genuine.

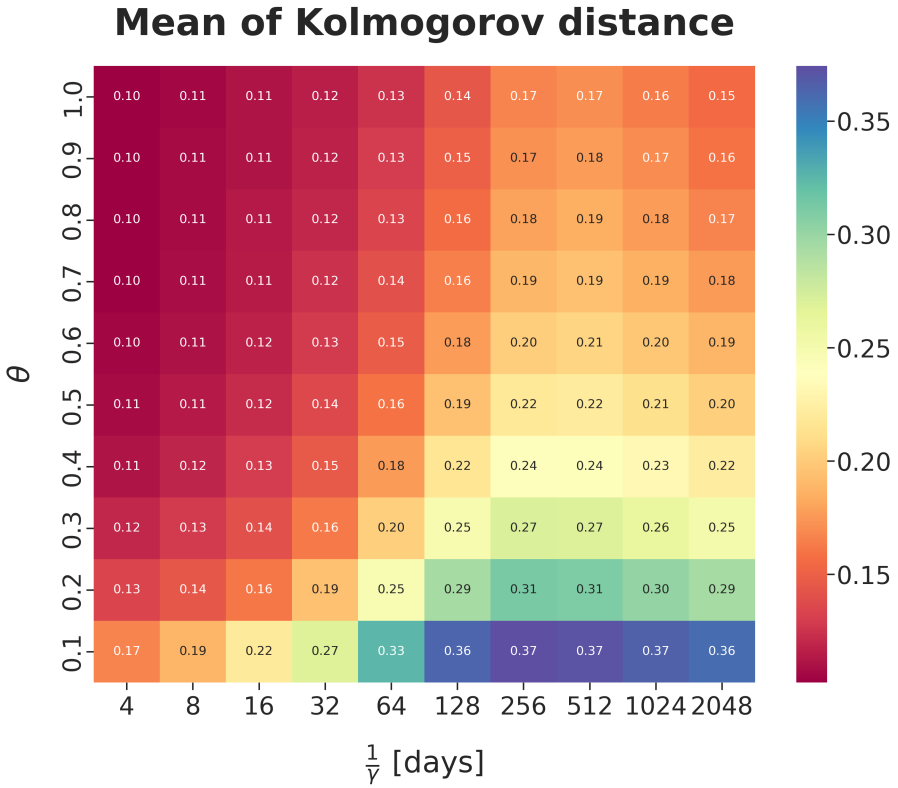

Figure 5.2: The Optimization Landscape (Heatmap). The color gradient represents the mean KS distance (lower values indicate better fit).

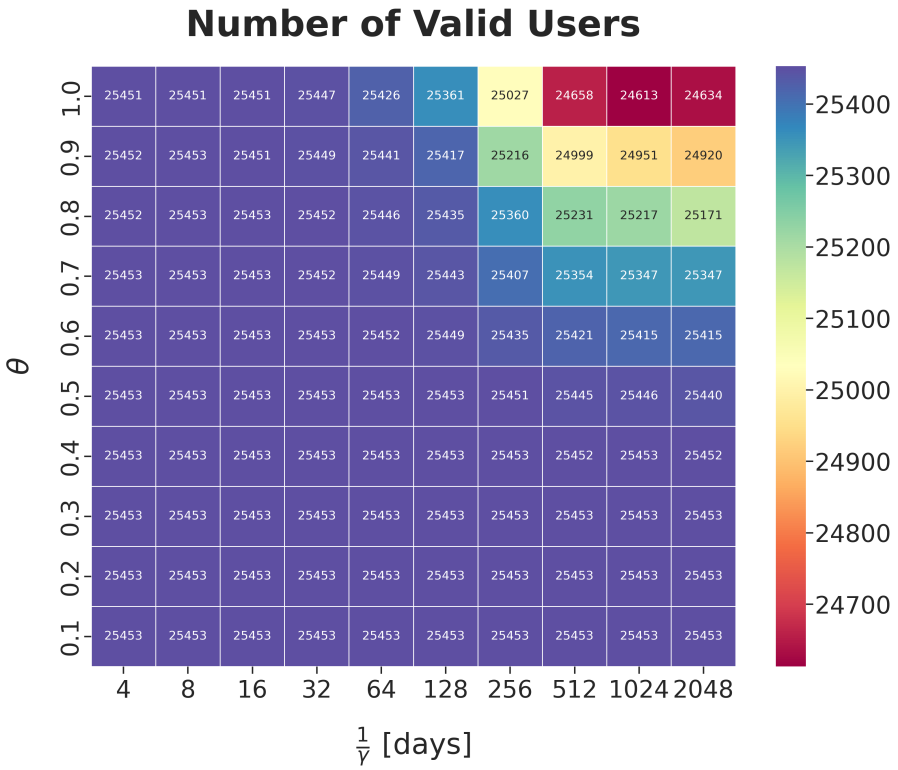

Figure 5.3: The number of valid users.

Beyond this empirical evidence, the selection of $\theta^* = 0.9$ is strongly justified by theoretical consistency with Theorem 1 of the reference framework. A value of $\theta \geq 1$ would imply a linear or super-linear Preferential Attachment mechanism. In the presence of multiplicative noise, such a regime is non-ergodic: the system would fail to reach a stationary distribution, leading to a scenario where popularity either diverges to infinity or concentrates entirely on a single node (monopoly).

The finding of $\theta^* = 0.9$ indicates that Bluesky operates in a *sub-linear* regime that is remarkably close to criticality. This specific tuning allows the platform to foster strong viral dynamics while mathematically guaranteeing the existence of a stable equilibrium, preventing the ecosystem from collapsing into a winner-takes-all state.

The Decay Rate (γ) and System Memory

The optimal decay rate corresponds to a time constant (or system memory) of:

$$\tau^* = 4 \text{ days} \tag{5.7}$$

This result is striking when compared to the high inertia observed in Facebook ($\tau_{FB} \approx 128$ days). A memory of just 4 days implies that Bluesky is an extremely volatile environment. Social capital depreciates rapidly; an influencer who stops posting effectively "disappears" from the collective attention field soon. This rapid decay implies that popularity on Bluesky is fleeting. Users are forced to remain constantly active to maintain their visibility; if they stop posting, their influence fades within days. This high volatility is likely a consequence of the chronological feed structure, which prioritizes new content over past engagement. The result is a highly dynamic environment: while it allows new users to gain visibility quickly, it also requires continuous effort from creators to avoid being forgotten.

5.2.3 Validation: Global Goodness-of-Fit

To validate the reliability of the calibrated model, we performed a comprehensive goodness-of-fit analysis on the entire user population. According to Assumption 2 of the theoretical framework, the normalized popularity jumps (Z) must follow a Lognormal distribution. We compared the fit of the Lognormal distribution against two alternative hypotheses: the Exponential distribution (memoryless) and the Power Law distribution (scale-free). The results, summarized in Figure 5.4, are clearly conclusive.

The analysis reveals that the Lognormal distribution provides the best statistical fit for 94.9% of the population.This validates the core assumption of the Stochastic Differential Equation: the interaction mechanism on Bluesky is governed by multiplicative noise. The rejection of the Exponential model (4.5%) confirms that engagement is not a random Poissonian process,while the negligible success of the Power Law (0.6%) for individual jumps indicates that a scale-free distribution is not

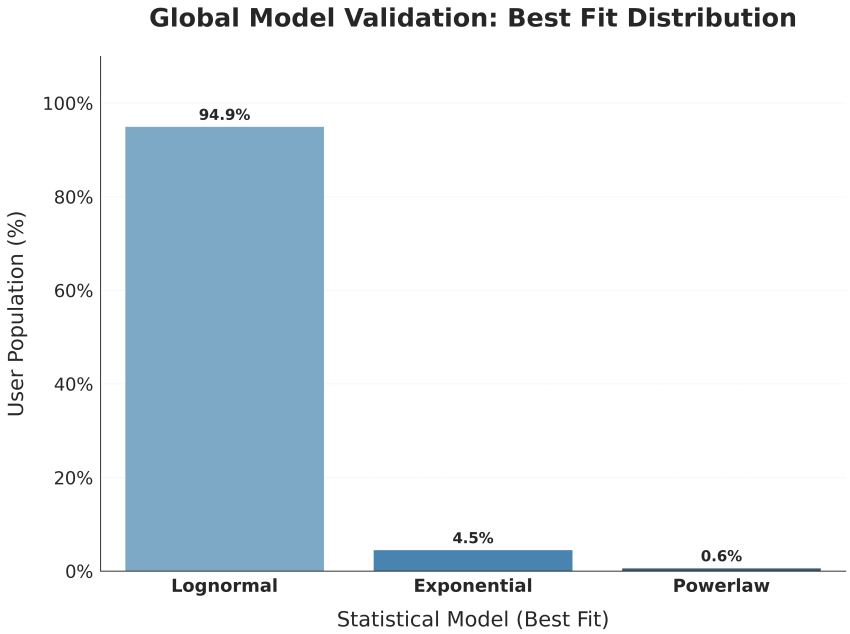

Figure 5.4: Global Validation of the Jump Process. .

the best fit for single interactions. This suggests that the heavy tails observed at the aggregate level may result from the cumulative process of the system's dynamics rather than from the statistical properties of individual jumps.

5.3 Social Stratification: The Two Regimes

The global analysis presented in the previous sections treats the Bluesky population as a homogeneous aggregate. However, a deeper investigation is required to determine whether the interplay between individual parameters (β_i, ϕ_i) and system dynamics (θ, γ) generates uniform outcomes or leads to distinct regimes of "Dominance" versus "Fair Play".

To verify whether the platform's rules depend on acquired status, we stratified the entire sample. This classification is not arbitrary but strictly aligns with the behavioral clusters previously identified through the K-Means unsupervised learning method presented in Section 4.2.1. Specifically, we focus on two distinct groups: the Influencers, corresponding to the High Engagement class, and the Wannabes, corresponding to the Medium Engagement class. Throughout the remainder of this analysis, these terms will be used exclusively to denote these algorithmically derived engagement profiles.

We repeated the Grid Search optimization procedure independently for each sub-population to verify if the macroscopic parameters governing the ecosystem shift as users climb the social ladder. The results, visualized in the heatmaps of Figure 5.5, reveal a striking consistency.

As shown in Figure 5.5, the optimization landscapes for both the Influencer and Wannabe classes mirror the global topology discussed in the previous section. Both groups converge towards a high viral capacity ($\theta \approx 0.9$) and a rapid decay rate. This

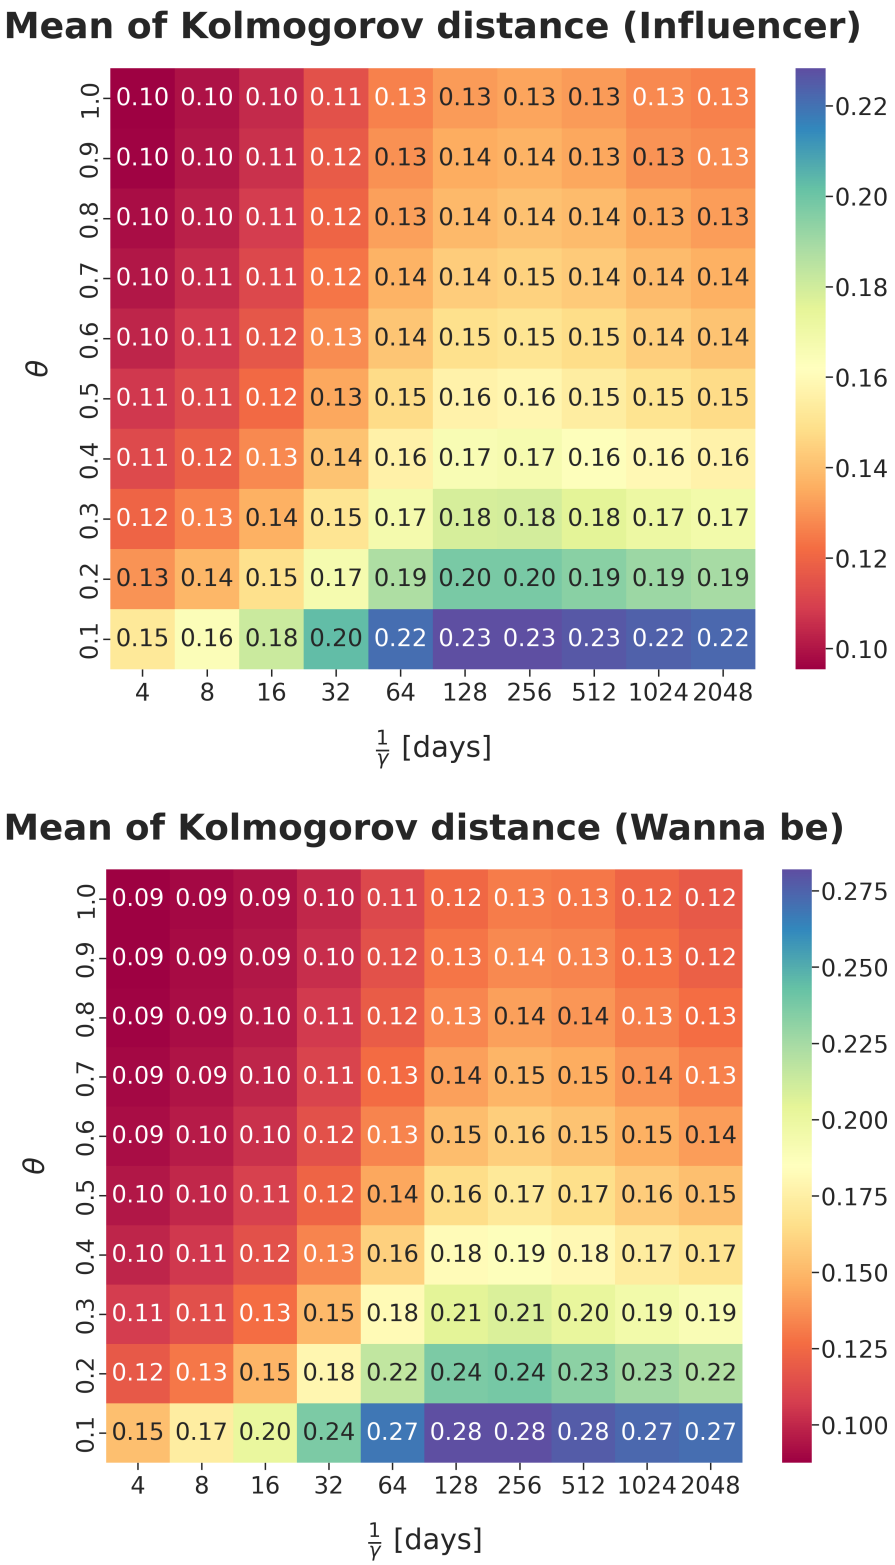

Figure 5.5: Comparison of the Optimization Landscapes: Influencer Class (top) and Wannabe Class (bottom). Both exhibit a global minimum in the same region of stability.

visual evidence suggests that the "rules of the game" on Bluesky are structural and democratic: the algorithm does not apply distinct underlying mechanisms to the

famous. A post by an influencer and a post by a Wannabe correspond to the same viral exponent θ; the massive difference in their actual performance is therefore not due to a change in the system's laws (θ, γ), but is entirely driven by their individual initial conditions $(X(t))$ and intrinsic merit (β_i), as we will explore in the breakdown of individual parameters.

5.3.1 Class-Specific Calibration

Adopting the Maximum Likelihood Estimation (MLE) methodology, we performed a separate calibration for the two groups. The results are summarized in Table 5.1.

Table 5.1: Comparison of structural parameters between the two regimes. The activity scaling parameter (ϕ) is intentionally omitted for the Wannabe class due to model misspecification in the sub-critical regime.

Parameter	**Influencer (Top 6%)**	**Wannabe (68-94%)**
Viral Capacity (θ)	0.90	0.90
Activity Scaling (ϕ_{tot})	0.0191	N/A
Activity Scaling ($\phi_{creative}$)	0.0173	N/A
Quality Variance (σ_β)	2.09	1.43
Burstiness (D)	12.85	10.32

The most striking finding concerns the **Viral Capacity (θ)**, which appears identical ($\theta \approx 0.9$) for both groups. Contrary to the hypothesis that the platform artificially amplifies the visibility of single posts by famous users, this result suggests that the *preferential attachment* mechanism at the individual content level is democratic: all else being equal, a post by a Wannabe possesses the same potential viral power as that of an Influencer.

However, a fundamental structural divergence emerges regarding the **Activity Scaling (ϕ)**. The parameter ϕ is an exponent of scale designed to capture non-linear network feedback (the 'Matthew Effect'). Estimating such an exponent rigorously requires the observable variable to span multiple orders of magnitude. The Wannabe class is, by definition, bounded within a narrow popularity domain ($X < X_c$, where X_c is the threshold for viral ignition).

In this sub-critical regime, forcing a power-law fit constitutes a severe model misspecification. It generates spurious, mathematically unreliable exponents driven by the boundaries of the domain (edge effects) rather than true network dynamics. Consequently, for the Wannabe class, the non-linear feedback loop ϕ is treated as inactive and mathematically negligible.

The true social fracture, therefore, is driven by the direct empirical metrics, specifically the **Merit Distribution (β)** and the **Temporal Strategy (D)**:

- **Quality Variance (σ_β):** Influencers exhibit higher variance ($\sigma \approx 2.09$) compared to Wannabes ($\sigma \approx 1.43$). In a lognormal distribution, a higher σ implies a heavier right tail. This indicates that while Wannabes produce content with consistent performance clustered around the mean, Influencers show a much

wider variability in content reception. Statistically, this higher variance grants the Influencers group a significantly larger probability of generating extreme positive outliers (viral posts), which are mathematically unlikely to occur in the narrower distribution of the Wannabes.

- **Temporal Strategy and Burstiness (D):** The analysis of the dispersion index ($D = \sigma_\lambda^2/\bar{\lambda}$) reveals that both classes drastically deviate from a random, memoryless Poisson process ($D = 1$). In fact, both groups exhibit an intensely "bursty" communication style ($D \gg 1$), indicating a systemic baseline adaptation to the platform's fast-paced environment. However, a behavioral divergence remains: the Elite class pushes this strategy further, operating at an even higher level of temporal concentration ($D = 12.85$ for Influencers vs. $D = 10.32$ for Wannabes).

  To ground this mathematically, the empirical activity distribution (visualized in the comparative boxplot, Figure 5.6) confirms the baseline divergence: Influencers maintain a higher typical output (median 3.34 posts/day) and a wider interquartile range compared to Wannabes (median 2.53 posts/day). Moving beyond daily averages, the inter-event time distribution explicitly highlights their clustered activity: the median wait time between posts is just 1.06 hours for the Elite (accelerating to 0.27 hours during peak bursts), compared to a 1.61-hour baseline for Wannabes.

  This statistical difference must be interpreted in the context of the system's short memory ($\tau^* = 4$ days), discovered during the global calibration (Section 5.2). Because collective attention fades so rapidly, a steady, low-frequency posting schedule is structurally inefficient. If a user waits too long between publications, the visibility generated by a post completely dissipates before the next one is published. This effectively resets their engagement to zero, making it mathematically impossible to accumulate social capital over time.

  The higher Burstiness of the Influencers ($D \approx 12.85$) compared to the Wannabes ($D \approx 10.32$) indicates an optimized adaptation to this algorithmic hostility. Rather than distributing their activity uniformly, Influencers cluster their publications into dense, high-frequency "bursts" or campaigns. This temporal strategy temporarily overwhelms the system's decay function (γ), stacking the visibility of multiple posts (V_i) within the short memory window. In contrast, the lower dispersion index of the Wannabes suggests a more organic and reactive use of the platform. While still exhibiting human-like irregularity, they lack the extreme, strategic synchronization of the Elite, leaving them more vulnerable to the rapid obsolescence of the timeline.

5.3.2 Stability Analysis and Growth Regimes

The structural difference between the two regimes emerges clearly when analyzing the stability condition derived in Theorem 1 of the Galante et Al. paper. The theorem

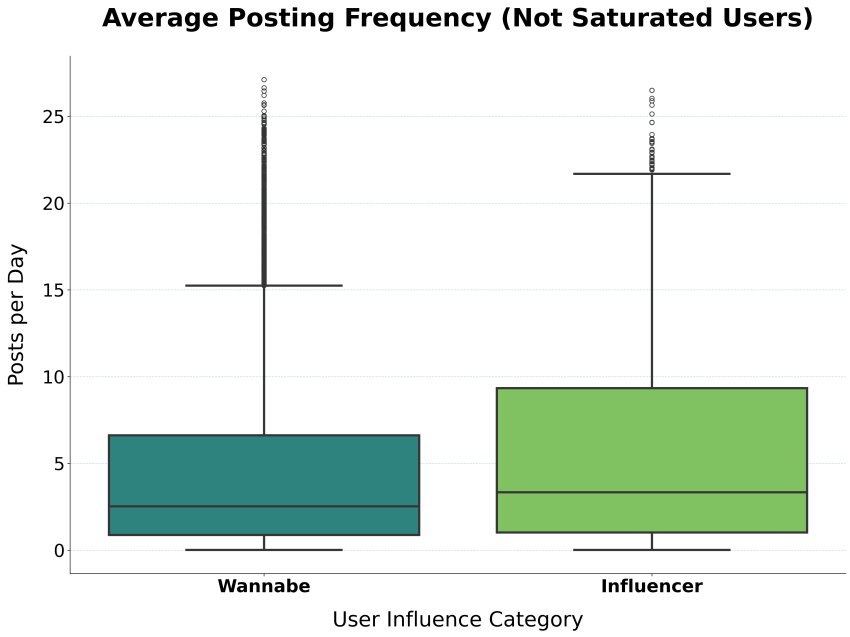

Figure 5.6: Comparative boxplot of average daily posting frequency between the High (Influencers) and Medium (Wannabes) engagement classes.

states that a user's growth potential is governed by the total exponent $\alpha = \theta + \phi$.

1. The Influencer Regime (Near-Critical Growth): For the dominant group, the sum of the parameters is:

$$\alpha_{inf} = \theta + \phi_{tot} \approx 0.9 + 0.0191 = 0.9191 \tag{5.8}$$

This value, very close to 1, places Influencers in a near-critical regime. The positive term $\phi \approx 0.02$ indicates a virtuous feedback loop: as popularity increases, the capacity to generate visibility increases slightly, allowing popularity to become self-sustaining. The resulting stationary distribution is Heavy-Tailed, permitting extreme excursions in popularity.

2. The Wannabe Regime (Baseline Sub-Criticality): For the majority of users, the situation is drastically different. Since the popularity $X(t)$ remains below the critical threshold (X_c) required to trigger autonomous network feedback, the non-linear behavioral scaling is treated as inactive (effectively omitting ϕ from the equation). Consequently, their growth potential is governed solely by the platform's baseline viral capacity:

$$\alpha_{wan} \approx \theta = 0.9 \tag{5.9}$$

While $\theta = 0.9$ permits isolated viral explosions, the absence of the additive ϕ term leaves the exponent strictly sub-critical ($\alpha_{wan} < 1$). Mathematically, this lack of behavioral feedback implies that the system exerts a dominant "restoring force" toward the mean. This defines a statistical Glass Ceiling [15]: a Wannabe might achieve a sudden spike in visibility, but without the self-sustaining behavioral feedback loop enjoyed by the Influencers, this attention is inevitably dampened by the platform's decay rate, preventing the long-term accumulation of social capital.

5.3.3 Drift and Restoration: The Role of Decay

While the static analysis of the exponent α suggests a structural limit to growth, the dynamic behavior of the system is best understood by examining the temporal evolution of popularity. Assuming the platform's attention allocation evolves via small, frequent increments, we transition from the discrete stochastic process to a continuous-time mean-field approximation. To formalize the competition between a user's effort and the platform's obsolescence, we derived the *Drift Equation*:

$$\frac{d\mathbb{E}[X]}{dt} \approx \underbrace{-\gamma \mathbb{E}[X]}_{\text{Decay Force}} + \underbrace{\lambda_0 \mathbb{E}[X]^{\theta+\phi}}_{\text{Growth Force}} \tag{5.10}$$

This differential equation reveals the instantaneous balance between the deterministic loss of attention (scaling linearly with the platform's obsolescence rate γ) and the regenerative capacity of new content, which scales as a power law of $\alpha = \theta + \phi$.

By setting Equation 5.10 to zero, we can identify the theoretical steady-state popularity, $\mathbb{E}[X]^*$, toward which a user naturally converges, a standard approach in the analysis of dynamical systems [**citare**]:

$$\mathbb{E}[X]^* = \left(\frac{\lambda_0}{\gamma}\right)^{\frac{1}{1-(\theta+\phi)}} \tag{5.11}$$

Equation 5.11 demonstrates that the equilibrium point is highly sensitive to the scaling exponent. Applying the calibrated parameters from Table 5.1, we observe two opposing dynamic regimes characterized by vastly different stability landscapes:

The Wannabe Trap (Decay Dominance): For the wanna-be class, the restorative force is heavily sub-linear since the growth relies entirely on θ without the compounding effect of ϕ ($\alpha \approx 0.9$). Because the denominator $1 - \theta$ remains relatively large compared to the Influencer regime, the resulting steady-state $\mathbb{E}[X]^*$ is low and acts as a strongly attractive stable fixed point. In simple terms, as soon as a Wannabe gains popularity through a stochastic spike, the system's physiological decay (γ) becomes disproportionately stronger than their baseline ability to push back. The platform mathematically forces the user back to their baseline equilibrium, making it nearly impossible to consolidate a viral spike into permanent fame.

The Influencer Advantage : For the Elite class, the growth exponent is near-unity ($\alpha \approx 0.92$). This implies that the regenerative force of new content grows almost in parallel with the decay, fundamentally altering the system's phase space. As the denominator $1 - (\theta + \phi)$ approaches zero, the equilibrium point $\mathbb{E}[X]^*$ shifts orders of magnitude higher. The system's friction is minimized, and the restoring force becomes marginally stable rather than strictly punitive. The attention lost to time is automatically replenished by the topological advantage and visibility of their established network. This allows Influencers to "float" at high popularity levels with a self-sustaining mechanism: the drift is essentially neutral, effectively removing the structural resistance that strictly bounds the Wannabes.

Chapter 6

Conlusion

6.1 Executive Summary of Findings

This thesis provided a comprehensive quantitative assessment of the popularity dynamics within the Bluesky ecosystem. To bypass the discoverability limitations of the decentralized architecture, we developed a custom multi-threaded crawler directly interfacing with the AT Protocol firehose. This empirical effort enabled the reconstruction of full 180-day historical timelines for a stratified sample of nearly 30,000 active users, yielding a longitudinal dataset of over 31 million distinct post. Leveraging this dataset, we modeled popularity not as a static metric, but as a continuous stochastic process $X(t)$ governed by opposing forces of content amplification and temporal decay. Specifically, we applied a Stochastic Mean-Field framework defined by the following Stochastic Differential Equation (SDE):

$$dX^{[i]}(t) = -\gamma X^{[i]}(t)dt + V_i(t)N_I^{[i]}(dt) + W_i(t)N_E^{[i]}(dt)$$

Through Maximum Likelihood Estimation (MLE) via Grid Search, we calibrated this model to isolate the global parameters governing the platform's attention economy. The empirical calibration revealed two structural properties:

- **System Memory and Decay:** The optimal decay rate γ corresponds to a characteristic system memory of $\tau^* = 4$ days. This reveals an ecosystem defined by extremely short-lived attention and high volatility, where social capital depreciates rapidly compared to established centralized platforms.
- **Viral Capacity and Ergodicity:** The viral exponent was empirically identified as $\theta^* = 0.9$. This demonstrates that the platform operates in a sub-linear regime remarkably close to criticality. This specific tuning allows for strong viral dynamics while mathematically guaranteeing the existence of a stable equilibrium, preventing the network from collapsing into a winner-takes-all monopoly.

Finally, the global goodness-of-fit validation confirmed that the endogenous popularity jumps $V_i(t)$ follow a lognormal distribution for 94.7% of the user population.

This validates the theoretical assumption that interaction mechanisms on the platform are governed by multiplicative noise rather than simple memoryless Poissonian processes.

6.2 The Volatility Paradox: Technical Freedom vs. Cognitive Scarcity

The Bluesky ecosystem is defined by a fundamental paradox between technological abundance and biological limits. From a technical perspective, the AT Protocol offers unprecedented freedom, empowering users with true "algorithmic choice". It allows the community to continuously generate a virtually infinite supply of custom feeds. However, this decentralized infrastructure clashes with the strict limits of human cognitive capacity. Even with a virtually infinite supply of algorithms, the aggregate cognitive bandwidth users can allocate cannot scale proportionally, maintaining attention as a fundamentally scarce resource.

This scarcity is empirically quantified by the "Zero-Like Phenomenon," where 63% of the developed algorithms (feeds) fail to overcome the initial discovery barrier, accruing at most a single like. This severe lack of visibility is the result of two combined factors: the practical difficulty users face in actually finding these new feeds, and the extremely rapid decay of collective attention, which resets the system's memory in just $\tau^* = 4$ days.

This rapid decay is striking when compared to the 128-day inertia typically seen on centralized platforms like Facebook. Because collective attention fades so quickly on Bluesky, creators are essentially forced to post continuously just to maintain their baseline visibility. In the end, our data points to a clear conclusion: while a decentralized architecture grants users more freedom to choose their algorithms, it cannot expand their cognitive limits or resolve the fundamental scarcity of human attention.

6.3 Social Stratification and mobility

Interestingly, our class-specific calibration reveals that the platform's core amplification mechanism treats users equally. The viral capacity ($\theta \approx 0.9$) is identical for both the established "Influencers" and the aspiring "Wannabes". The algorithm does not automatically favor the elite at the individual post level. Instead, the true social fracture is driven by behavioral strategies and initial network topology.

Influencers maintain their dominance by operating in a near-critical growth regime. They adapt to the rapid attention decay by clustering their content into high-frequency campaigns, effectively stacking visibility before the system's short memory resets their progress. In contrast, Wannabes interact more organically and remain trapped in a sub-critical state.

This dynamic effectively maintains an algorithmic glass ceiling. A decentralized infrastructure allows emerging actors to go viral occasionally, but the fierce metabolic

cost of attention prevents them from converting these transient spikes into permanent, long-term social mobility.

6.4 Final Considerations and Future Directions

Bluesky represents a successful experiment in creating a decentralized digital environment. However, our findings demonstrate that technical freedom does not automatically guarantee social equality. The algorithm itself is neutral, but human tendencies and established networks still generate biases. Since attention drops rapidly, these natural inequalities grow without specific algorithmic interventions. Ultimately, despite using an open protocol, the platform mirrors the strict social hierarchies of centralized networks.

To move beyond the Algorithmic Glass Ceiling and foster genuine social mobility, future research and protocol development should explore two main directions:

- **Algorithmic Bias Correction:** Standard recommendation systems tend to amplify "homophily" (the tendency of users to interact only with similar, already-connected profiles). This creates an echo chamber that locks out new creators. Developers of custom Feed Generators should build specific mathematical corrections into their algorithms. By actively pushing content outside of a user's usual network, these corrected algorithms can break the glass ceiling and offer fair visibility to emerging accounts.
- **Anti-Decay Mechanisms:** The current system memory of just 4 days structurally punishes users who cannot maintain continuous, high-frequency activity. Future research should explore the implementation of "anti-decay" algorithms capable of identifying and preserving the visibility of high-quality content over longer periods. This approach would decouple algorithmic success from sheer posting volume, allowing for a slower, yet truly meritocratic, accumulation of social capital.

Bibliography

[1] Martin Kleppmann, Paul Frazee, Jake Gold, Jay Graber, Daniel Holmgren, Devin Ivy, Jeromy Johnson, Bryan Newbold, and Jaz Volpert. "Bluesky and the at protocol: Usable decentralized social media". In: *Proceedings of the ACM Conext-2024 Workshop on the Decentralization of the Internet*. 2024, pp. 1–7 (cit. on pp. 1, 5, 7).

[2] Dorian Quelle and Alexandre Bovet. "Bluesky: Network topology, polarization, and algorithmic curation". In: *PloS one* 20.2 (2025), e0318034 (cit. on pp. 1, 5, 12, 13).

[3] Franco Galante, Chiara Ravazzi, Luca Vassio, Michele Garetto, and Emilio Leonardi. "Dominance or Fair Play in Social Networks? A Model of Influencer Popularity Dynamics". In: *International Conference on Advances in Social Networks Analysis and Mining*. Springer. 2025, pp. 306–321 (cit. on pp. 1, 2, 4, 10, 29, 40).

[4] Ozgur Can Seckin, Filipi Nascimento Silva, Bao Tran Truong, Sangyeon Kim, Fan Huang, Nick Liu, Alessandro Flammini, and Filippo Menczer. *The Rise of Bluesky*. 2025. arXiv: `2504.12902 [cs.SI]`. URL: `https://arxiv.org/abs/2504.12902` (cit. on pp. 5, 6).

[5] Ali Salloum, Dorian Quelle, Letizia Iannucci, Alexandre Bovet, and Mikko Kivelä. "Politics and polarization on Bluesky". In: *arXiv preprint arXiv:2506.03443* (2025) (cit. on pp. 5, 12, 13).

[6] Dorian Quelle, Frederic Denker, Prashant Garg, and Alexandre Bovet. *Simple contagion drives population-scale platform migration*. 2026. arXiv: `2505.24801 [cs.SI]`. URL: `https://arxiv.org/abs/2505.24801` (cit. on pp. 6, 12).

[7] Lucio La Cava, Luca Maria Aiello, and Andrea Tagarelli. "Drivers of social influence in the Twitter migration to Mastodon". In: *Scientific Reports* 13.1 (2023), p. 21626 (cit. on p. 6).

[8] Leonhard Balduf, Saidu Sokoto, Andrea Baronchelli, Ignacio Castro, Michał Król, Gareth Tyson, George Pavlou, Björn Scheuermann, and Onur Ascigil. "Bootstrapping social networks: Lessons from Bluesky starter packs". In: *Proceedings of the International AAAI Conference on Web and Social Media*. Vol. 19. 2025, pp. 178–192 (cit. on pp. 6, 7).

[9] Lilian Weng, Alessandro Flammini, Alessandro Vespignani, and Fillipo Menczer. "Competition among memes in a world with limited attention". In: *Scientific reports* 2.1 (2012), p. 335 (cit. on p. 7).

[10] Maxi Heitmayer. "The second wave of attention economics. Attention as a universal symbolic currency on social media and beyond". In: *Interacting with Computers* 37.1 (2025), pp. 18–29 (cit. on pp. 7, 8).

[11] Fang Wu and Bernardo A Huberman. "Novelty and collective attention". In: *Proceedings of the National Academy of Sciences* 104.45 (2007), pp. 17599–17601 (cit. on p. 8).

[12] Jacob Ratkiewicz, Santo Fortunato, Alessandro Flammini, Filippo Menczer, and Alessandro Vespignani. "Characterizing and modeling the dynamics of online popularity". In: *Physical review letters* 105.15 (2010), p. 158701 (cit. on pp. 8, 9).

[13] Philipp Lorenz-Spreen, Bjarke Mørch Mønsted, Philipp Hövel, and Sune Lehmann. "Accelerating dynamics of collective attention". In: *Nature communications* 10.1 (2019), p. 1759 (cit. on p. 8).

[14] Albert-Laszlo Barabasi. "The origin of bursts and heavy tails in human dynamics". In: *Nature* 435.7039 (2005), pp. 207–211 (cit. on p. 8).

[15] Ana-Andreea Stoica, Christopher Riederer, and Augustin Chaintreau. "Algorithmic glass ceiling in social networks: The effects of social recommendations on network diversity". In: *Proceedings of the 2018 World Wide Web Conference*. 2018, pp. 923–932 (cit. on pp. 9, 12–14, 48).

[16] Marian-Andrei Rizoiu, Lexing Xie, Scott Sanner, Manuel Cebrian, Honglin Yu, and Pascal Van Hentenryck. "Expecting to be hip: Hawkes intensity processes for social media popularity". In: *Proceedings of the 26th international conference on world wide web*. 2017, pp. 735–744 (cit. on p. 9).

[17] Yuwei Zhu and Paolo Barucca. "Parsimonious Hawkes Processes for temporal networks modelling". In: *arXiv preprint arXiv:2501.17720* (2025) (cit. on p. 10).

[18] Matjaž Perc. "The Matthew effect in empirical data". In: *Journal of The Royal Society Interface* 11.98 (2014) (cit. on p. 11).

[19] Albert-László Barabási and Réka Albert. "Emergence of scaling in random networks". In: *science* 286.5439 (1999), pp. 509–512 (cit. on p. 12).

[20] Réka Albert and Albert-László Barabási. "Statistical mechanics of complex networks". In: *Rev. Mod. Phys.* 74 (1 Jan. 2002), pp. 47–97. DOI: 10.1103/RevModPhys.74.47. URL: https://link.aps.org/doi/10.1103/RevModPhys.74.47 (cit. on p. 12).

[21] Andrea Failla and Giulio Rossetti. ""I'm in the bluesky tonight": insights from a year worth of social data". In: *PloS one* 19.11 (2024), e0310330 (cit. on p. 30).

%Dedications