

High Tension Beliefs in High Ambiguity Information Environments: A Framework for Analyzing COVID-19 Origin Theories on Twitter/X

Samantha Clark¹, Scott Leo Renshaw¹, and Kathleen M. Carley¹

Carnegie Mellon University
Pittsburgh, PA
srenshaw@cs.cmu.edu

Abstract. The COVID-19 pandemic created conditions of high importance and ambiguity that, according to Allport and Postman’s rumor theory, facilitate alternative narrative formation. This study analyzes COVID-19 laboratory origin theories on Twitter during January-June 2020, applying Stark and Bainbridge’s concept of “tension” to examine belief communities that maintained opposition to mainstream scientific consensus. We analyzed tweets using BotHunter and BotBuster for automated account detection and BEND maneuver analysis for influence tactics. Our research addresses three key questions: What proportion of our dataset contains laboratory origin theories versus mainstream content? How do different bot detection algorithms identify automated activity in these contested spaces? Do automated accounts maintain consistent engagement over time? Through established frameworks of rumor transmission and social tension, we examine how technological affordances and social dynamics interact during crisis periods. This analysis contributes to understanding how alternative explanatory frameworks emerge and propagate during uncertain events, while establishing a methodological framework for future research to systematically analyze the content dynamics and actor behaviors within high tension belief communities operating in high rumor propensity environments. This framework provides a foundation for investigating similar phenomena across different crisis contexts and contested information domains.

Keywords: COVID-19 · Communication Strategies · High Tension Beliefs · Information Ambiguity · Online Social Media · Bot Behavior

1 Introduction

Allport and Postman’s classic model posits that a rumor’s strength equals importance multiplied by ambiguity – where rumor circulation increases with both the thematic significance of events and uncertainty surrounding them [1]. The COVID-19 Pandemic, and particularly the first few months, created one of the

most significant contexts for information transmission and rumoring in recent history.

This paper examines the “Chinese lab creation theory” – the belief that COVID-19 was intentionally created in a laboratory in Wuhan, China. During our study period (January-June 2020), no conclusive reports confirmed laboratory creation as the origin, representing a period of significant uncertainty and ambiguity that fostered alternative explanatory frameworks.

Rather than focusing on the theory’s veracity, we employ the concept of “high tension beliefs” from research on unconventional religious movements [10]. High tension groups are described as those “in a high state of tension with their sociocultural environment,” while mainstream groups remain in “low tension” with societal norms.

2 Literature Review

2.1 Theoretical Framework and COVID-19 Context

The COVID-19 pandemic created ideal conditions for what Allport and Postman identified as high rumor propensity environments, where rumor strength equals importance multiplied by ambiguity [1]. Beginning with the WHO’s Public Health Emergency declaration on January 30, 2020 [13], and escalating through the U.S. national emergency [12] and lockdown guidelines [11], the pandemic combined maximum personal significance with fundamental uncertainty about viral origins.

During this period, laboratory origin theories represented what Stark and Bainbridge termed “high tension beliefs” – positions that do not necessarily cohere with prevailing cultural and evidence-based consensus [10]. While institutional sources emphasized zoonotic origins, lab creation theories gained traction online throughout early 2020, punctuated by events like: President Trump’s April 30 claims of having seen evidence (although providing no concrete data at the time) [5] and Tucker Carlson’s May allegations about gain-of-function research funding [7]. These high tension beliefs operated in direct opposition to mainstream scientific consensus during our study period (January-June 2020), when no conclusive reports confirmed laboratory creation as the origin.

This framework allows examination of belief community positioning without predetermined assumptions about claim validity, focusing instead on information dynamics during periods of scientific and social uncertainty. The framework recognizes that tension levels shift over time as evidence emerges, and that high tension beliefs may serve important social functions during uncertainty.

2.2 Empirical Research on COVID-19 Information Dynamics

Recent research has examined contested information propagation during the pandemic. Understanding conspiracy theories requires recognizing their psychological and social functions, particularly during periods of uncertainty and social

threat [6]. Moffitt et al. analyzed COVID-19 tweets achieving over 90% accuracy in classifying conspiracy-labeled content, revealing that high tension belief communities relied on more concentrated information sources (5,485 unique URL domains versus 10,178 for mainstream content) and utilized hashtags more frequently [8]. Bot detection revealed similar automated participation across content types (33% in non-conspiracy versus 36% in conspiracy-labeled content), indicating widespread rather than targeted automated engagement [8, 14].

Building on this foundation, our analysis employs Sentence Transformers for content classification and dual bot detection algorithms (BotHunter and BotBuster) to examine automated participation patterns. We apply the BEND maneuver framework to analyze strategic communication tactics, examining how actors navigate contested discourse spaces during periods of maximum uncertainty [4].

3 Data

3.1 Data Collection

Our analysis draws upon a comprehensive dataset of COVID-19-related Twitter communications collected via the Twitter V1 API, encompassing the critical early pandemic period from January through June 2020. This temporal scope captures the initial phase of global uncertainty when competing explanatory frameworks for viral origins were emerging within online discourse communities. The dataset comprises 14,714 tweets extracted from a larger COVID-19 corpus using MongoDB and filtered through strategically selected keywords designed to capture high tension beliefs regarding laboratory origins: “chinese lab,” “WuhanVirus,” “ChineseVirus,” “ChinaVirus,” and “ChinaLiedPeopleDied.”

These filtering parameters were intentionally constructed to identify content that positioned itself in opposition to mainstream zoonotic origin narratives prevalent during this period. The inclusion of terms that explicitly implicate Chinese laboratory involvement, combined with hashtags that suggest deliberate concealment or deception, allows us to examine how high tension belief communities constructed and disseminated alternative causal explanations during a period of maximum ambiguity regarding viral origins.

4 Methodology

Our analytical approach integrates multiple computational methods to examine high tension belief propagation during the high ambiguity and uncertainty around crisis. The analytical pipeline begins with comprehensive data preprocessing using MongoDB for tweet extraction, followed by emotional feature extraction using Netmapper to generate CUES (Cognitive, Understand, Emotional, Social) scores for subsequent BEND maneuver analysis. Bot detection employs both BotHunter and BotBuster algorithms: BotHunter examines semantic features, user metadata, temporal patterns, and interaction behaviors [2], while

BotBuster uses machine learning models trained on diverse datasets to predict bot behavior when traditional indicators are incomplete [9]. This dual-algorithm approach enables robust bot detection assessment across different computational frameworks.

Following bot classification, BEND (Build, Engage, Narrow, Dismiss) maneuver analysis uses ORA software to identify influence tactics within high tension content, building upon emotional analysis to reveal how actors shape discourse in contested information environments.

4.1 BEND Maneuver Framework

The BEND framework provides systematic categorization of influence operations within social media environments [3], examining tactical methods for shaping public discourse rather than focusing solely on content veracity.

Table 1. BEND Maneuvers Classification Schema

	Community	Narrative
Constructive	Back, Build, Bridge, Boost	Excite, Explain, Engage, Enhance
Destructive	Negate, Neutralize, Narrow, Neglect	Dismay, Distort, Dismiss, Distract

The taxonomy distinguishes between community-oriented maneuvers focusing on network relationships and narrative-oriented maneuvers targeting message content and framing. The constructive-destructive dimension examines whether influence tactics build alternative explanatory systems or undermine institutional credibility, providing insights into strategic orientations within contested information spaces while avoiding normative judgments about content validity.

5 Results

5.1 Tweets and Platform Affordances

Table 2. Tweet Breakdown by Month

Month	Number of Tweets	Number of Agents	Hashtags	URLs
Jan 2020	2284	11207	502	170
Feb 2020	2290	9501	522	182
Mar 2020	2618	2496	513	183
Apr 2020	2605	2595	980	265
May 2020	2858	2575	1124	337
June 2020	2059	2075	106	12




Fig. 1. Platform affordances over time showing the number of tweets, hashtags, and URLs from January to June 2020. Dotted vertical lines indicate March and May as key time periods.

Table 2 reveals several notable patterns in engagement dynamics across the study period. The substantial decrease in unique agents from January (11,207) to March (2,496) suggests a consolidation of active participants within high tension belief communities, potentially indicating the formation of more concentrated discussion networks as initial uncertainty gave way to more polarized positions. Simultaneously, the dramatic increase in hashtag usage from April onward, peaking at 1,124 in May, suggests intensified efforts at message amplification and community building during this period. The precipitous decline in June across all metrics likely reflects shifting public attention toward other significant events, including emerging civil rights demonstrations following George Floyd’s death.

5.2 Conspiracy Content and Bot Activity Patterns


“Conspiracy”-labeled tweets fluctuated significantly over the six-month study period, peaking in March 2020 with 315 tweets compared to just 15 in June. This surge coincided with the early pandemic lockdowns and heightened public uncertainty about COVID-19 origins.

Both BotHunter and BotBuster algorithms revealed concerning trends in automated account activity, though with notably different detection patterns. BotHunter identified a troubling shift in the conversation landscape: while human participation plummeted by 84% (from 7,493 to 1,188 accounts), bot activity declined by only 72% (from 2,721 to 765 accounts). This meant bot activity relative to human participation nearly doubled, rising from 26% to 39% between January and June.

Month	Conspiracy Labeled	Non-conspiracy Labeled
Jan 2020	278	3,722
Feb 2020	157	1,842
Mar 2020	315	1,684
Apr 2020	91	1,908
May 2020	219	1,780
June 2020	15	1,984

Table 3. Lab Creation Conspiracy vs Non-conspiracy Tweets by Month

BotBuster painted a more conservative picture, detecting bot activity ranging from 15.6% to 14.7% relative to humans, with more similar decline rates for both groups (humans 80%, bots 81.6%). Despite these methodological differences, both algorithms confirmed that bots maintained stronger engagement in lab leak discussions even as overall public interest waned.

**Fig. 2.** Bot versus human activity detection results comparing BotHunter and BotBuster methodologies from January to June 2020. Dotted vertical lines indicate March and May as key time periods.

The conspiracy content analysis revealed March as the peak month for bot-generated conspiracy tweets across both detection methods. BotHunter identified 64 conspiracy-tweeting bots in March, while BotBuster found 41—both representing the highest monthly totals in the dataset.

Perhaps most telling was the episodic nature of bot activity. The vast majority of automated accounts did not maintain consistent activity across all six months, suggesting these weren't part of sustained influence campaigns but rather event-driven responses to shifting public discourse. Both algorithms iden-




Fig. 3. Bot-generated conspiracy tweets over time comparing detection results from BotHunter and BotBuster methodologies from January to June 2020. Dotted vertical lines indicate March and May as key time periods.

tified different account sets, with only one shared account (“ma”) across the entire study period, indicating algorithm-specific detection thresholds. Peak persistence occurred during January-February and March-April transitions, corresponding to heightened pandemic attention periods.

Table 4. Cross-Month Bot Account Persistence: BotHunter vs BotBuster Comparison

Month	BotHunter	BotBuster
Jan-Feb 2020	p1, r1, and f1	gen, oc, and yh
Feb-Mar 2020	f1	no matches
Mar-Apr 2020	mk, s1, ds1, and rr1	tt
Apr-May 2020	su and cq	no matches
May-June 2020	ma and mz	ma

5.3 Communication Strategies: BEND Maneuvers

The analysis of communication strategies revealed consistent patterns in how different tactics were deployed throughout the study period. “Engage” appeared as a top strategy every single month, making it the most persistent maneuver at 100% frequency. “Narrow” ranked second, appearing in the top three strategies in four out of six months (66.67%), and notably dominated whenever it appeared.

March marked a strategic shift with the first appearance of “Explain” maneuvers, which continued into April—possibly reflecting attempts to provide

Month	Top Three BEND Maneuvers
Jan 2020	Narrow, Engage, and Distract
Feb 2020	Engage, Distract, and Narrow
Mar 2020	Engage, Distract, and Explain
Apr 2020	Narrow, Engage, and Explain
May 2020	Narrow, Engage, and Boost
June 2020	Engage, Build, and Bridge

Table 5. Top Three BEND Maneuvers by Month

alternative explanations during peak conspiracy activity. May stood out as the only month featuring “Boost” strategies, while June showed a distinct pivot toward community-building tactics with “Build” and “Bridge” maneuvers appearing for the first time. This June shift toward inclusive strategies coincided with the overall decline in conversation volume, suggesting efforts to re-engage audiences as interest waned.

The strategic evolution from narrow, divisive tactics in early months to more explanatory and community-building approaches later in the period suggests adaptive communication strategies responding to changing audience engagement and platform dynamics.

6 Discussion

The temporal distribution of Chinese laboratory creation theory content exhibits characteristic rise-and-fall patterns linked to competing contemporary events. While overall tweet volume remained consistent (averaging 2,452 tweets), engagement patterns revealed a notable paradox: stable participation metrics alongside declining unique agents, indicating intensified participation among a concentrated user base. This concentration effect peaked in May 2020, with hashtag usage reaching 1,124 instances—twice the six-month average—before collapsing to 106 in June. This precipitous decline aligns with the emergence of Black Lives Matter protests following George Floyd’s death, suggesting competing high-salience events redirected social media attention away from pandemic origin discussions.

Comparative bot detection analysis reveals that methodological choice significantly influences prevalence estimates. BotHunter identified concerning increases in bot activity relative to participation (26% to 39% from January to June), while BotBuster demonstrated more conservative and stable patterns (15.6% to 14.7%). Both algorithms converged on peak automated activity during March 2020, coinciding with maximum conspiracy content and U.S. lockdown implementation, indicating event-driven rather than sustained coordinated campaigns.

BEND maneuver analysis reveals regular strategic patterns throughout the period. “Engage” appeared consistently across all months, while “Narrow” dominated when present (January, April, May). The notable shift occurred in June, where strategies pivoted to “Build” and “Bridge”—representing compensatory

efforts to rebuild and expand conversation participation as overall interest declined following external events.

These findings suggest that automated activity in high tension belief spaces demonstrates responsiveness to external events and strategic adaptation to maintain influence despite declining engagement. The differential decline rates between human and bot participation, combined with evolving communication strategies, indicate sophisticated approaches to maintaining disproportionate influence within shrinking conversation spaces.

7 Limitations and Future Directions

Our findings are constrained by several methodological limitations. The keyword filtering strategy introduces systematic bias by targeting adversarial framings of Chinese involvement, potentially excluding neutral scientific discussions and limiting analysis to politically charged discourse rather than comprehensive origin speculation. Additionally, temporal distance between data collection and analysis has rendered portions of the dataset inaccessible due to account suspensions and platform policy changes, constraining longitudinal user behavior tracking.

Future research should expand the temporal scope beyond six months to examine trend persistence across different crisis contexts and validate classification methodology through alternative conspiracy content identification approaches and cross-validation frameworks. Extended longitudinal analysis would enhance understanding of belief community evolution patterns during high-tension periods.

8 Conclusion

This analysis of Chinese laboratory creation theory propagation demonstrates how high tension beliefs emerge and evolve during periods of crisis uncertainty, validating Allport and Postman’s foundational insights regarding rumor strength as a function of importance and ambiguity. The temporal convergence of peak conspiracy content, automated activity, and U.S. lockdown implementation in March 2020 illustrates how high tension belief communities exploit periods of maximum ambiguity, while the dramatic engagement decline following George Floyd’s death in May demonstrates how competing high-salience events can redirect social attention away from contested origin narratives.

The differential decline rates between human and bot participation (84% vs 72% respectively), combined with evolving BEND maneuver patterns from “Narrow” and “Distract” tactics to “Build” and “Bridge” approaches, reveals how automated systems maintain disproportionate influence within shrinking conversation spaces by adapting to shifting information landscapes. These findings suggest that high tension belief propagation operates according to predictable crisis dynamics: flourishing during periods of institutional uncertainty while remaining vulnerable to attention capture by competing events that offer greater immediacy or emotional resonance than abstract origin theories.

References

- [1] Gordon W. Allport and Leo Postman. *The Psychology of Rumor*. Henry Holt and Company, New York, 1947.
- [2] David Beskow and Kathleen Carley. Bot-hunter: A tiered approach to detecting & characterizing automated activity on twitter. 07 2018.
- [3] Janice Blane. *Social-Cyber Maneuvers for Analyzing Online Influence Operations*. Ph.d. thesis, Carnegie Mellon University, Pittsburgh, PA, 2023. Accessed: 2025-04-11.
- [4] Kathleen M. Carley. Social cybersecurity: An emerging science. *Computational and Mathematical Organization Theory*, 26(4):365–381, 2020.
- [5] CNN. Trump contradicts us intel community by claiming he’s seen evidence coronavirus originated in chinese lab. News Report, April 30 2020.
- [6] Karen M. Douglas, Joseph E. Uscinski, Robbie M. Sutton, Aleksandra Cichocka, Turkay Nefes, Chee Siang Ang, and Farzin Deravi. Understanding conspiracy theories. *Political Psychology*, 40:3–35, 2019.
- [7] Fox News. Tucker carlson: Why the media, scientific community and who won’t investigate covid origins. Opinion Commentary, March 31 2021.
- [8] J. D. Moffitt, Catherine King, and Kathleen M. Carley. Hunting conspiracy theories during the covid-19 pandemic. *Social Media + Society*, 7(3):1–12, 2021. Accessed: 2025-04-11.
- [9] Lynnette Hui Xian Ng and Kathleen M. Carley. Botbuster: Multi-platform bot detection using a mixture of experts. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 17, pages 686–697. AAAI Press, 2023. Accessed: 2025-04-11.
- [10] Rodney Stark and William Sims Bainbridge. *A Theory of Religion*. Peter Lang, New York, 1987.
- [11] The White House. 15 days to slow the spread. Government Guidelines, March 16 2020.
- [12] The White House. Proclamation on declaring a national emergency concerning the novel coronavirus disease (covid-19) outbreak. Presidential Proclamation, March 13 2020.
- [13] World Health Organization. Statement on the second meeting of the international health regulations (2005) emergency committee regarding the outbreak of novel coronavirus (2019-ncov). WHO Press Release, January 30 2020. WHO declares Public Health Emergency of International Concern.
- [14] Kai-Cheng Yang, Francesco Pierri, Pik-Mai Hui, David Axelrod, Christopher Torres-Lugo, John Bryden, and Filippo Menczer. The covid-19 infodemic: Twitter versus facebook. *Social Media + Society*, 7(3):1–12, 2021. Accessed: 2025-04-11.