The Impact of Offspring Hashtags on Semantic Polarization in Online Social Movements: Evidence from the Indian Farmers' Protest

Rohan Singh Leekha

Thesis submitted to the Faculty of the

Virginia Polytechnic Institute and State University

in partial fulfillment of the requirements for the degree of

Master of Science

in

Computer Science and Applications

Eugenia Rho, Chair

Ismini Lourentzou

Edward Fox

May 3, 2023

Blacksburg, Virginia

Keywords: offspring hashtags, online social movements, discourse analysis, polarization, Indian farmers' protest, social media

The Impact of Offspring Hashtags on Semantic Polarization in Online Social Movements: Evidence from the Indian Farmers' Protest

Rohan Singh Leekha

(ABSTRACT)

In this work, we investigate the role of offspring hashtags on the semantic polarization of online discourse between the protest and counter-protest communities over time through the lens of the 2021 farmers' protest in India. Offspring hashtags are those that first appear alongside their more widely known "parent" hashtag (e.g., #WhyIDidntReport and #YesAllWomen are offspring hashtags that first co-appeared alongside their more famous and mainstream parent hashtag, #MeToo). The prominence of parent hashtags and their visible role in facilitating modern day protests have dominated scholarly efforts in understanding the socio-technical influence of social movement hashtags. By contrast, scholarship on the impact of the lesser-known offspring hashtags is rare and typically examined through the lens of its primary parent tag. Our work aims to address this gap. In this research, we examine how the protest and counter-protest communities use offspring hashtags in their tweets to discuss and frame farmers - the key social group at the center of the farmers' protest (RQ1). Our findings reveal that both protests and counter-protests use offspring hashtags in a manner that further polarizes rather than bridges perspectives on core issues - focusing on themes that malign the other side (RQ2). We then measure and track how the semantic polarization in the use of the term "farmer" by the protest vs. counter-protest communities who use offspring hashtags evolves over time in relation to key protest events (RQ3). Finally, to empirically test and demonstrate whether and how the volume of offspring hashtags throughout the protest period influences semantic polarization trends between the protest and counter-protest discussion of farmers, we create a series of time-series models for causal

inference. We use Granger-causality to test whether and how fluctuations in the volume of offspring hashtags significantly predict how the protest and counter-protest communities semantically diverge in how they discuss farmers over time (RQ4). By analyzing offspring hashtags, this work provides a detailed understanding of the nuanced themes and narratives that may be lost under parent hashtags, but significantly influence online discourse between the protest and counter-protest communities.

The Impact of Offspring Hashtags on Semantic Polarization in Online Social Movements: Evidence from the Indian Farmers' Protest

Rohan Singh Leekha

(GENERAL AUDIENCE ABSTRACT)

In this study, we explore how offspring hashtags impact online conversations between people supporting and opposing the 2021 farmers' protest in India. Offspring hashtags are less popular hashtags that first appear with a more famous "parent" hashtag (for example, #WhyIDidntReport and #YesAllWomen alongside #MeToo). While researchers have extensively studied parent hashtags, the influence of offspring hashtags remains less explored. Our research looks at how protests and counter-protests use offspring hashtags to talk about farmers, who are at the center of the Indian farmers' protest. We found that both groups use offspring hashtags in a way that increases polarization rather than fostering understanding between opposing sides. This often leads to discussions that focus on attacking the other group. We also analyzed how the polarization in conversations about farmers evolved over time, in relation to key protest events and the use of offspring hashtags. To see if the number of offspring hashtags used during the protest affected polarization trends, we used statistical models and a method called Granger-causality. Our findings show that fluctuations in offspring hashtag volume significantly predict how protesters and counter-protesters diverge in their discussions about farmers over time. By examining offspring hashtags, we gain a deeper understanding of the subtle themes and stories that may be overlooked when focusing only on parent hashtags but play a crucial role in shaping online conversations between opposing groups.

Dedication

To

my mother, Neena Singh Leekha

and

father, Bobby Leekha

for their unconditional love, support, and encouragement

“Sometimes resistance is not a matter of choice. Not all battles are fought for victory. Some are fought for simply to tell the world that someone was there on the battlefield”. – Ravish Kumar (Ramon Magsaysay Award, 2019, Winner)

Acknowledgments

I express my sincere gratitude to my advisor Dr. Eugenia Rho, who has been a source of guidance, encouragement, and inspiration throughout my academic journey. I also extend my appreciation to my committee members and the departmental chair for their mentorship and guidance. I would also like to acknowledge my lab mates (Xiaohan, Buse, Kaike, Anisha and Uma), who have been a source of support throughout my time in the program. I have learned so much from them, both professionally and personally. I am grateful for their friendship and support. Finally, I would like to thank my parents for their unconditional love, support, and encouragement throughout my journey as a graduate student.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Whether it is about human rights [50], war [4], politics [85], or the environment [108], social media hashtags over the last decade have played a significant role in mobilizing protests and magnifying critical social issues, bringing them to the forefront of global conversations through online networks [75]. The democratizing power of online hashtags has been widely cited by scholars in Human-Computer Interaction (HCI) and Computer-Supported Cooperative Work (CSCW) who highlight the affordance of hashtags to scale conversations [7], expand audience reach [98], and amplify marginalized voices [35]. Yet, the impact of social media hashtags has also been associated with the rapid ascent of extreme polarization. Recent studies empirically link hashtags with the spread of partisan echo chambers [9, 112], political mis-/disinformation [81, 98], and hateful narratives [24, 32] across protest and counter-protest communities [36, 85]. One such example is the farmers' protest in India.

1.1 Background of the Farmers' Protest in India

With more than 250 million participants, the farmers' protest in India has been referred to as one of the largest protests in modern history [30]. From September 2020 to November 2021, millions of citizens in India rallied against the government to repeal three agricultural laws that would eliminate existing regulatory support for farmers by privatizing India's agricultural sector [30]. In response to the government's passing of the farm bills, hundreds of millions of farmers and their families gathered in a massive nationwide strike, coalescing on- and offline under the #FarmersProtest hashtag to voice concern against the government's decision seen as a threat to their livelihood [30, 63, 95]. Contesting this view, a counter-protest community quickly emerged online, responding to the #FarmersProtest by using the same hashtag to dilute opposition narratives and amplify pro-government views [79, 86] – a phenomenon known as hashtag-hijacking [117]. In no time, the #FarmersProtest became the center of media and public attention, generating widespread reactions from both sides of the issue on social media [86, 89, 101]. Soon after the emergence of the #FarmersProtest hashtag [95], the polarized split between the protest and counter-protest communities quickly paved the way for an outburst of hyperpartisan news coverage [79]. Sensationalized media framing [107], fake news [25], and extremist opinions [25, 79] about the protest rapidly proliferated through the #FarmersProtests hashtag on social network sites. Thousands of new tweets, TikTok, Instagram, and Facebook posts about the protest appeared every day tagged with the #FarmersProtest hashtag [89, 101]. As a result, the hashtag #FarmersProtest became the second most trending hashtag of the year 2021 [3] with journalists [63], celebrities [79], and politicians [79] from all over the world using the hashtag to offer their two-cents amid public rallies and physical altercations involving citizens from both sides of the protests.

The powerful affordance of social media hashtags [57, 109] was once again demonstrated through the rapid growth and virality of the #FarmersProtest hashtag across online conversations. However, the countless number of "offspring" hashtags that emerged alongside their "parent" tag, #FarmersProtest, have not received much attention. In this work, we define offspring hashtags as those that first co-appear with their more well-known and mainstream parent hashtag. Both the protest and counter-protest communities created their own offspring hashtags under the umbrella of the parent hashtag, #FarmersProtest. Table 1.1 showcases the top 20 offspring hashtags from the protest and counter-protest communities based on the descending frequency in which they appeared in our data. For example, along with the parent tag #FarmersProtest, the protest community used offspring hashtags, such as #IStandWithFarmers, #KisaanEktaZindabad, and #NoMoreBJP to express their support for the farmers and their opposition to the current government. In contrast, the counter-protest community used #FakeFarmers, #Rioters, and #ModiWithFarmers with the primary tag #FarmersProtest in their tweets to delegitimize the protest or claim that the farmers were being misled.

1.2 Analyzing the Indian Farmers' Protest Over Other Online Social Movements

The Indian farmers' protest offers a unique lens to examine a large-scale, predominantly economically driven movement in a non-western context, focusing on the impact of policy change on the livelihood of millions of small-scale farmers [11]. This multi-faith movement, led by religious minorities in India, showcases the complexities of navigating a socio-political landscape where majoritarianism frequently remains unchallenged, providing a distinct perspective on the motivations, strategies, and interactions between protesters, counter-protesters, and the government, as compared to movements like #BlackLivesMatter [7] and #MeToo [37]. In addition, the Indian farmers' protest highlights the role of digital activism in connecting geographically dispersed communities and garnering international support [30, 64, 79], showcasing the power of hashtags and offspring hashtags in shaping the discourse and amplifying diverse voices. The movement's unique characteristics, such as the interplay between religious minorities and government dynamics [11], as well as the cultural and social factors influencing the protest [48], make it an invaluable case study for understanding the intricacies and nuances of online social movements across various contexts and issues.

Pro-Farmer Offspring Hashtags	**Anti-Farmer Offspring Hashtags**
#IStandWithFarmers (1.62%)	#FakeFarmers (3.36%)
#SaveFarmers (1.58%)	#brokers (2.49%)
#FarmerSuicide (1.54%)	#ModiHaiToMumkinHai (2.44%)
#chakkajam (1.44%)	#KanganaRanuat (2.14%)
#DeathOfDemocracy (1.01%)	#bestpm (1.70%)
#NoFarmersNoFood (0.81%)	#KhalistanBehindViolence (1.67%)
#NoMoreBJP (0.79%)	#secessionism (1.67%)
#lathicharge (0.72%)	#separatist (1.53%)
#MediaCircus (0.69%)	#khalistan (1.50%)
#KisanMajdoorEktaZindabaad (0.66%)	#FakeFarmersProtest (1.50%)
#BJPdestroysDemocracy (0.66%)	#RahulPriyankaFakeDrama (1.47%)
#SpeakUpForFarmers (0.65%)	#FakeFarmersProtest (1.41%)
#ISupportBharatBandh (0.64%)	#Congress_Cheap_Politics (1.31%)
#KisaanFightsForRights (0.62%)	#noRoadJam (1.27%)
#kisaanvirodhinarendramodi (0.58%)	#KhalistaniTerrorists (1.05%)
#GodiMedia (0.56%)	#Communists (1.07%)
#noonecaresaboutfarmer (0.55%)	#Tabligh (0.90%)
#FarmersDilliChalo (0.55%)	#BoycottFood (0.87%)
#KisaanAndolan (0.53%)	#ShaheenBagh2 (0.80%)

Table 1.1: Top 20 Offspring Hashtags from the Protest and Counter-Protest Communities and their Percentage of Occurrence in the Dataset.

1.3 Motivation of Research

In this work, we investigate the role of such offspring hashtags on the semantic polarization[1] of public discourse between the protest and counter-protest communities through the lens of the 2021 farmers' protest in India. By "we", I specifically refer to myself and my lab member, who helped with data validation (Chapter 3, Section 3.2) for this study.

The prominence of parent hashtags and their visible role in facilitating or polarizing public discourse and protests have dominated scholarly efforts in understanding the socio-technical influence of social movement hashtags [52, 96]. By contrast, scholarship on the impact of lesser-known offspring hashtags in HCI and CSCW is rare and, typically examined through the lens of its primary parent tag [82]. Our work aims to address this gap by asking the

[1]We define semantic polarization as how protests and counter-protests semantically diverge across time in their contextual use of the keyword "farmer" through the course of the Indian farmers' protest.

following research questions (RQs).

RQ1: How do offspring hashtags emerge from the parent hashtag, #FarmersProtest? What is the pattern of emergence and growth of offspring hashtags in terms of speed and volume?

In RQ1, we aim to understand how offspring hashtags emerge from the parent tag, #FarmersProtest. What is the temporal and volumetric pattern of the emergence of offspring hashtags from the parent hashtag, #FarmersProtest, and how do they differ between the protest and counter-protest communities? By asking RQ1, we seek to understand the timing and extent to which the protest and counter-protest communities create and use offspring hashtags within the discursive networks that discuss the farmers' protest. Understanding the growth and diffusion of offspring hashtags allows us to understand their mechanism and potential for mobilization and polarization around the farmers' protest.

RQ2A: What kind of frames do protesters and counter-protesters employ to portray farmers in their discussion?

Prior research in HCI/CSCW [6, 57, 72] has shown that marginalized social groups are often subject to greater scrutiny during and after online protests [46] and are more likely to experience negative outcomes, such as violence [97], social exclusion [46], and economic disadvantages [43, 62], as a result of polarization. Therefore, in our examination of offspring hashtags and their role in semantically polarizing online social movements, we first seek to understand how farmers, one of India's largest socioeconomically marginalized demographic populations, are framed across the protest and counter-protest discussions. Framing refers to the way in which information is presented to influence how individuals understand and interpret a particular issue or event [105]. It involves selecting certain aspects of a topic and highlighting them over others in order to shape how individuals perceive and interpret the issue [105]. Different framing patterns can lead to different interpretations and understanding of the same issue. Therefore, framing can impact how individuals view the issue, what actions they may take in response to it, and what solutions they might propose. Hence, in RQ2A, we identify the different frames through which the protest and counter-protest communities discuss and depict farmers in the context of the Indian farmers' protest.

RQ2B: How do protesters and counter-protesters use offspring hashtags in their framing and discussion of farmers?

To better understand the role that offspring hashtags play in the framing of farmers in the protest and counter-protest discourse, we examined how the two groups used offspring hashtags in their discussion of farmers. Specifically, *how* do protesters and counter-protesters use offspring hashtags to construct and contest identities [61, 100] that surround their emerging perceptions of farmers throughout the online social movement? To answer RQ2B, we used qualitative analysis to delve into the different rhetorical strategies in which the two groups used offspring hashtags to frame their discussion of farmers.

RQ3: How does the semantic polarization in the usage of the term "farmer" evolve over time among protesters and counter-protesters who use offspring hashtags?

To capture how offspring hashtags contribute to the linguistic polarization between the protest and the counter-protest discourse at-scale, we calculated how semantic polarization between the two groups' discussion of the keyword "farmer" evolves over time. RQ2 shows that the protest and counter-protest communities who use offspring hashtags drastically differ in their framing and portrayal of farmers. Do these communities also *semantically* diverge in how they talk about farmers involved in the protest? In RQ3, we aim to capture the semantic polarization in how the two opposing groups of the farmers' protest (protest and counter-protest communities) use the term "farmer" in tweets that contain offspring hashtags throughout the entire protest period. How polarized are the two groups in the manner in which they discuss a key social group at the center of one of the largest social movements in modern history? How does this semantic polarization *evolve over time* in relation to key protest events? To answer RQ3, we capture semantic polarization by measuring the diachronic shifts [31, 76] in how the protest and counter-protest communities diverge in their use of the word "farmer" across tweets that contain offspring hashtags throughout the entire protest period (16 months from September 9, 2020 to November 11, 2021). By doing so, we show how semantic polarization evolves over time across key protest events as well as the volume of offspring hashtags associated with the protest.

RQ4: How does the use of offspring hashtags in discussions related to farmers by protest and counter-protest communities affect semantic polarization trends?

RQ3 shows that semantic polarization in which the protest vs. counter-protest communities talk about farmers closely evolves and fluctuates with key protest events as well as the volume of offspring hashtags across the protest period. Hence, in RQ4, we test the significance of this relationship. We ask whether the patterns in the volume of offspring hashtags throughout the protest period are significantly predictive of semantic polarization trends in how the protest and counter-protest communities discuss and portray farmers. To answer this question, we test Hypothesis 1 (H1) and Hypothesis 2 (H2) using Granger-causality analysis. H1 and H2 allow us to empirically test whether and how offspring hashtags drive semantic polarization in the protest and counter-protest communities' discussion of farmers.

H1: *The daily number of offspring hashtags significantly Granger-causes semantic polarization between the protests and counter-protest communities' discussion of farmers.*

H2: *The daily number of offspring hashtags that newly emerge each day significantly Granger-causes semantic polarization between the protests and counter-protest communities' discussion of farmers.*

1.4 Contributions

Our study is one of the first to provide empirical evidence of the significant impact of offspring hashtags on the semantic polarization of online discussions about a major social group during a social movement. As shown in the case of the Indian farmers' protest, offspring hashtags emerged almost immediately and rapidly, just fifteen minutes after the first appearance of the dominant and better-known parent hashtag. Less than a week into the protest (marked by the first appearance of the parent tag, #FarmersProtest), 1,202 offspring hashtags were circulating throughout the networked discourse across the protest and counter-protest communities, accounting for a total of 3.5% of the entire volume of offspring hashtags that emerged from the whole protest period of 16 months (RQ1). Offspring hashtags emerge early on, quickly, and in large volumes at the very beginning stages of an online social movement, meaning its impact on the polarization of protest conversations is significant as we demonstrate in this work. Our findings reveal that both protests and counter-protests use offspring hashtags in a manner that further polarizes rather than bridges perspectives on core issues - focusing on themes that malign the other side (RQ2). Counter-protesters use offspring hashtags to portray farmers as "rapists", while protesters depict farmers as victims of an authoritarian government similar to that of "Hitler". Further, the semantic polarization in how these two groups talk about farmers evolves closely with key protest events as well as the volume of offspring hashtags (RQ3). In fact, the fluctuating volume of offspring hashtags throughout the protest period is a significant and reliable predictor of semantic polarization patterns in how the protests and counter-protests diverge in their framing and discussion of farmers (RQ4). Key findings of our study have significant implications for understanding the impact of offspring hashtags created by protest and counter-protest groups on polarizing networked discussions between the two communities. Our research sheds light on how online protest hashtags can potentially contribute to the spread of extremist views, further exacerbating the division between opposing protest camps rather than bridging it.

Chapter 2

Review of Literature

2.1 How Parent (and Offspring) Hashtags Power Social Media Protests

Prior research in HCI and CSCW have demonstrated early on that hashtags are effective tools for organizing online social movements [66]. For example, the well-known #MeToo [56] and #BlackLivesMatter [121] highlight the importance of hashtags in stimulating critical conversations around social media protests [66]. Whether calling people to action, challenging existing narratives, or sharing personal experiences, people participating in online social movements are able to amplify their messages far beyond their original audience through the affordances of hashtags [82, 99]. People use hashtags to assert social identities, promote group causes, or bring attention to perspectives that emerge throughout the protest [36, 47, 75, 110]. The subsequent conversation, fueled by hashtags, can rapidly gain traction and momentum during a protest, as millions of users who are interlinked through a common hashtag engage in discourse across the network.

However, it is not only through the well-known and often viral *parent* hashtags like #MeToo and #BlackLivesMatter, but also the lesser known *offspring* hashtags, through which users engage in protest discourse [15, 23, 78]. For example, during the #MeToo movement, offspring hashtags like #WhyIDidntReport and #YesAllWomen "spawned" or first co-appeared with the dominant parent tag #MeToo. Similarly, #saytheirnames and #justiceforgeorgefloyd appeared for the first time through posts hashtagged with the trending parent tag

#BlackLivesMatter.

Empirical studies have demonstrated that offspring hashtags, which are created in conjunction with parent hashtags, reveal specific linguistic patterns and themes that drive discourse within protest and counter-protest communities [7, 97]. Unlike their parent counterparts, offspring hashtags are typically more specific and nuanced [91, 110], serving as a form of meta-commentary [122] or introducing perspectives that challenge the main narrative represented by the parent tag [110]. However, despite their significance, parent hashtags receive more attention from researchers and media when it comes to understanding their role in polarizing online social movements [36, 91].

This study aims to examine discourse beyond the broader umbrella of the parent tag, #FarmersProtest. We investigate how protest and counter-protest communities use offspring hashtags to contest, negotiate, and assert various frames in their portrayal of farmers throughout the Indian farmers' protest. By analyzing the use and evolution of offspring hashtags over time, we aim to understand their impact on semantic polarization between the protest and counter-protest communities.

2.2 The Role of Political Hashtags in Polarizing Public Discourse

While parent hashtags originating out of social movements can quickly mobilize communities and create a sense of collective identity around specific issues [110], they can also contribute to the creation of online partisan echo chambers [83]. These echo chambers can harden or confirm ideological biases [36] across like-minded individuals [36, 110]. For example, researchers conducted a large-scale online experiment to determine whether the presence of hashtags (#BlackLivesMatter and #MeToo) in social media news articles affects the quality of discourse around social topics related to race and gender [96, 97]. The results showed that compared to the control group, those who were shown hashtags in their news posts perceived the news content as less socially important and reported less motivation to know more about social issues related to the post [97]. Furthermore, people who viewed news posts with political hashtags wrote comments with significantly more words associated with anger, disgust, and fear than those who viewed the same news article without the hashtag [97].

Such studies however, solely focus on the parent hashtags instead of lesser-known offspring hashtags. Meanwhile, studies that examine offspring hashtags typically focus on those that have reached a critical mass or gained significant audience exposure (e.g., #BluesLivesMatter or #AllLivesMatter), rather than those that emerge throughout the protest. Consequently, we know little about the impact of newly spawned offspring hashtags, especially those that never reach a critical mass, in polarizing online social movements. In particular, we do not know whether and how offspring hashtags contribute to discourse polarization, and whether they do so to a greater or lesser extent than parent hashtags. Our research aims to address these gaps in the literature by investigating the impact of offspring hashtags on polarizing discourse between protest and counter-protest communities that emerged during the Indian farmers' protest.

2.3 Capturing Semantic Polarization Between Protests and Counter-Protests Over Time

In previous research, the quantification of linguistic polarization has been treated as a static measure derived from a single longitudinal data dump [26, 87, 104]. As a result, such methods commonly adopted in CSCW and HCI research on online social movements tend to provide an averaged snapshot of polarization across hashtagged conversations [5, 68, 69].

These approaches are computationally lightweight, but only offer an aggregated measure of polarization at a particular point in time. Hence, they overlook the temporal dynamics that are essential for understanding how polarization evolves over time [2, 103].

Recent advances in Natural Language Processing (NLP) research using diachronic embeddings have provided a better understanding of how the contextual meaning of a word changes over time [67, 76]. Scholars have shown that diachronic shifts, or changes in how a word is used across time [67, 76], reflect dynamical shifts in societal and cultural consensus around important topics [67]. For example, diachronic shifts using contextual word embeddings with temporal features have been computed to demonstrate how certain keywords ("racism", "health care", "climate change", "police", etc.) have been discussed in broadcast media language over the last ten years, and how these shifts have contributed to peaks in political polarization [31]. Understanding diachronic shifts in the contextual use of a keyword is important as it enables us to comprehend how language and meaning are constructed by society [67]. By observing how the meaning of a keyword evolves over time, we can better grasp the social dynamics that influence the use of language in a specific context. These semantic shifts can also capture pivotal moments in which protest discourse takes a significant turn [49, 116].

In this study, we aim to gain a better understanding of how farmers, who represent the primary social group at the center of the Indian farmers' protest, are discussed across protest and counter-protest communities using offspring hashtags *throughout* the course of the protest. To overcome prior methodological challenges, we adopt methodological intuitions from NLP research on diachronic embeddings. Specifically, we use the framework proposed by [31] to measure semantic polarization by capturing the diachronic shifts in how protest and counter-protest communities diverge in their use of the word "farmer" over the course of the Indian farmer's protest across posts that contain offspring hashtags.

Chapter 3

Data

To ensure the data used in our study accurately reflects the discourse surrounding the Indian farmers' protest, we systematically collected tweets from September 9, 2020, to November 11, 2021, using the parent hashtag, "#FarmersProtest". We based our data collection on two main criteria:

- Tweets containing the parent hashtag and at least one offspring hashtag.
- Tweets originating from India (based on the Twitter geo-tag) and written in English.

3.1 Data Pre-processing

We implemented the following pre-processing steps for the tweets in our dataset before further analysis:

- All emoticons and media links were removed.
- Removed the "RT" keyword from the beginning of retweets.
- Removed punctuation marks and extra spaces.
- Replaced digits with the word equivalent.
- Converted all characters to lowercase letters.

3.2 Data Validation

We followed a three-step process to classify tweets as protest or counter-protest, using offspring hashtags:

- **Step 1:** The thesis author and one additional lab member manually evaluated a randomized sample of 800 tweets from the corpus. Through this evaluation, we identified 1000 offspring hashtags - 500 from the protest community and 500 from the counter-protest community. Both raters independently evaluated the same sample of offspring hashtags. Following the independent evaluation, inter-rater reliability (IRR) was assessed using Cohen's Kappa (k) statistic [77], which quantifies the degree of agreement between two raters, adjusting for chance agreement. The kappa coefficient (k) quantifies IRR by comparing the actual agreement among raters (observed agreement) with the level of agreement that could happen randomly (expected chance agreement). The Cohen's Kappa statistic was determined to be 0.85, indicating a high level of agreement between the two raters. All discrepancies between the two raters were resolved through multiple rounds of discussions and additional examination of tweets that contained the offspring hashtags in question.
- **Step 2:** Next, we used the list of 500 protest and 500 counter-protest offspring hashtags from Step 1 to further classify additional tweets from our corpus. If a tweet contained an offspring hashtag from the list of offspring hashtags from the protest community, we labeled it as a protest tweet. Similarly, if a tweet included an offspring hashtag from the counter-protest list, we labeled it as a counter-protest tweet. If any newly discovered tweets contained unseen offspring hashtags, they were added to the respective lists of protest or counter-protest offspring hashtags.
- **Step 3:** We re-scanned the entire dataset, using our updated list of protest and counter-protest offspring hashtags to label previously unseen tweets.

We repeated Step 3 until no new offspring hashtags or tweets were left to be found.

Our final dataset consisted of a total of 241,531 tweets (239,581 protest and 1,680 counter-protest tweets) (see Table 3.1). There were a total of 36,930 offspring hashtags (34,696 protest and 2,234 counter-protest offspring hashtags) in our data (see Table 3.2). Among scholarship [60, 80, 86, 90, 118] that analyze tweets from the Indian farmers' protest, our dataset was the largest (the next largest dataset includes 169,000 tweets [60]).

Descriptive Statistics	**Protest**	**Counter-Protest**	**Total**
Total number of tweets	239,851	1,680	241,531
Mean number of words per tweet	13	21	14
Median number of words per tweet	8	20	8
25th percentile of words per tweet	3	10	4
75th percentile of words per tweet	21	32	21

Table 3.1: Descriptive Statistics of Tweets Identified as Protest and Counter-Protest.

Descriptive Statistics	**Protest**	**Counter-Protest**	**Total**
Total number of offspring hashtags (unique)	34,696	2,234	36,930
Mean number of offspring hashtags per tweet	2	4	2
Median number of offspring hashtags per tweet	1	3	1
25th percentile number of offspring hashtags per tweet	0	2	0
75th percentile number of offspring hashtags per tweet	2	5	2

Table 3.2: Descriptive Statistics of Offspring Hashtags Identified as Protest and Counter-Protest.

Chapter 4

Methods

4.1 RQ1: How do Offspring Hashtags Emerge from the Parent Hashtag, #FarmersProtest?

After obtaining the dataset, we begin the network construction process. This process starts with initializing an empty graph that includes a single node representing the parent hashtag (#FarmersProtest). The parent hashtag acts as the central node in the graph. To prevent duplication in the network, we maintain a set of unique hashtags, ensuring that each hashtag is included only once [114].

The network's nodes represent the parent hashtag and its offspring hashtags. The central hashtag node serves as the network's focal point, while the offspring hashtag nodes connect to the central node based on their co-occurrence with the central hashtag in the dataset. Each node can be depicted as an object containing the hashtag text [13].

We establish the network's edges based on the relationship between the parent hashtag and its offspring hashtags. An edge is created between the parent hashtag node and an offspring hashtag node if the two hashtags co-occur in the same tweet [51].

To create the network structure, we iterate over each tweet in the dataset and extract the hashtags present in the text. If the central hashtag appears in the entry, we establish connections between the central hashtag and the other hashtags within the same entry. We then add each unique hashtag as a node and connect it to the central hashtag with an edge.

To organize the network structure, we use the force-directed layout algorithm [17]. This algorithm positions nodes in the network by simulating a system of forces acting upon each node and edge, resulting in a visually appealing and easy-to-interpret representation [17].

Finally, to visualize the constructed network, we use the NetworkX [74] visualization library to draw the nodes representing each hashtag with a specific size and color. We also depict the edges connecting the nodes to showcase the relationships between the central hashtag and its offspring hashtags [74]. We label each node in the network with its corresponding hashtag.

4.2 RQ2A: What Kind of Frames do Protesters and Counter-Protesters Use to Portray Farmers in their Discussion?

Identifying contextually similar word pairs is a widely adopted methodological paradigm in understanding how a topic or an issue is framed [19, 98]. Therefore, to answer RQ2A, we first identified keywords that are contextually closest to how the protest and counter-protest groups, respectively use the word "farmer" in their posts that contained offspring hashtags. This allows us to understand how the two groups frame and portray farmers in their discussion. For this task, we first used the Sentence-BERT (SBERT) [1] model to learn and generate contextual word embeddings for all words (each of which is a 384×1 tensor) [93], including the token "farmer" in every occurrence in both the protest and counter-protest corpora of tweets.

SBERT [1] is a large language model used to generate dense vector representations, or embeddings, for sentences [1]. The methodology used by SBERT involves fine-tuning the BERT model [28] using a Siamese or Triplet network structure, which allows it to learn fixed-size vector representations for input sentences in a supervised manner. We use the SBERT model to generate embeddings of words in a sentence by feeding the sentence through the model, with each word in the sentence mapped to its corresponding embedding in the output vector [1, 113]. SBERT takes into consideration the context and meaning of words within a sentence [20], which allows us to capture the connotation and nuance of each instance of the word "farmer" as expressed in a tweet by the protest and the counter-protest communities. In this research, we use the pre-trained "all-MiniLM-L12-v2" SBERT architecture with 12-layer, 384-hidden size dimension, 12-heads, 21M Transformer parameters and 96M embedding parameters [38].

Once we derived word embeddings for all words in our corpus, we used Facebook AI Similarity Search (FAISS) [59] to identify words that were most contextually similar to the token "farmer". FAISS is an open-source library developed by Facebook AI Research for performing efficient similarity search and clustering of high-dimensional vectors [59]. It normalizes high dimensional embeddings and then applies dimensionality reduction before organizing them into an index for efficient search [58]. FAISS uses cosine similarity to measure the semantic similarity between the query vector and indexed vectors, enabling fast and accurate search results [73]. We used FAISS [111] to calculate cosine values to quantify the contextual similarity between the word embeddings of all the words in our corpus (indexed vectors) and the word embedding for the token "farmer" (query vector). This allowed us to identify the top 20 words that were contextually most similar to the word "farmer" in each of the protest and counter-protest corpora of tweets based on cosine values. Cosine values for word embedding pairs range from -1, indicating the most contextually different, to 1, signifying the most contextually similar.

4.3 RQ2B: How do Protesters and Counter-Protesters Use Offspring Hashtags in their Framing and Discussion of Farmers?

Next, to understand how the protest and counter-protest communities use offspring hashtags in their framing of farmers, we employed qualitative analysis. In RQ2A, we identified the top 40 most contextually similar words to the token, “farmer” as used by the protesters and the counter-protesters in our data (20 words for each group). We then filtered all tweets that included at least one occurrence of one of the 40 tokens in our data, which resulted in a total of 2000 tweets (1000 tweets from the protest and 1000 tweets from the counter-protest corpora). We then proceeded to manually examine a randomly selected sample of tweets for each token occurrence, leading us to examine a total of 1150 tweets (600 from protests and 550 from counter-protests). We first adopted an inductive coding [42] process to allow themes to emerge from our initial examination of tweets that included the top 40 tokens. We then iteratively discussed comparing and organizing codes to identify the dominant themes in which these tokens were discussed in context. This enabled us to categorize the predictive tokens into key thematic groups based on their contextual use in sentences [40]. We resolved any disagreements through multiple rounds of discussions before confirming a final set of thematic categories. Finally, we used the same sample to conduct discourse analysis [40] to gain a deeper understanding of how the protest and counter-protest communities used offspring hashtags to convey their underlying intent in their framing of farmers. Discourse analysis is a qualitative research method that entails discerning patterns, relationships, and underlying values within textual data to gain insight into the intended meaning behind the text [40]. We, therefore, used this method to examine the language of the protest and counter-protest communities in our data, as well as how each community used offspring hashtags to convey their underlying intent and message in their tweets.

4.4 RQ3: How Does the Semantic Polarization in the Usage of the Token "farmer" Evolve Over Time Among Protesters and Counter-Protesters Who Use Offspring Hashtags?

To answer RQ3, we use the hidden layers from the pre-trained Bidirectional Encoder Representations from Transformers (BERT) model to extract the diachronic shift in the meaning of the token "farmer" in protest and counter-protest communities.
BERT is a pre-trained neural network model for natural language processing [28]. It is used to create context-aware word embeddings, which are numerical representations of words or phrases that convey their meaning and context. BERT achieves this by bi-directionally processing text, enabling it to comprehend text from both directions and consider the entire context of a sentence or paragraph [28]. BERT processes text bidirectionally, allowing it to understand the relationships between words more effectively. Unlike other machine learning techniques such as Word2Vec, term frequency-inverse document frequency (tf-idf), etc. [28, 31], which only consider local context by using a fixed window around a word, BERT captures the entire context of a word in a sentence or paragraph[1] [64]. This enables BERT to produce richer and more accurate word embeddings that are better suited for a wide range of NLP tasks such as embeddings generation and text classification [28, 31, 64]. We use a bert-base-uncased model, which has 12 layers, 16 self-attention heads, a hidden size of 768 and 110M parameters. It is pre-trained on the BookCorpus, a dataset consisting of 11,038 unpublished books and English Wikipedia [28].

Our method follows the methodology outlined in [31], with the exception of the input keyword being "farmer" instead of the keyword used in the original study. First, we use a pre-trained

[1]BERT has a length limit on the number of tokens considered, so if the paragraph is very long, it only captures part of the full context.

BERT [28] model to capture embeddings of the token "farmer". Then, for each occurrence of the token "farmer", we calculate the cosine distance between the contextual embeddings generated in the protest and counter-protest communities. We capture the change in word meaning for each day, where both protest and counter-protest communities use the token "farmer" to quantify the semantic polarization of discourse between the two groups. Our calculation of semantic polarization of how protest and counter-protest communities diverge in the contextual understanding of the token "farmer" is as follows.

Given the representation of a word embedding for the token "farmer" for protest tweets:

$$P_t = \{w_i, w_{i+1}, w_{i+2}\} \tag{4.1}$$

and the representation of a word embedding for the token "farmer" for counter-protest tweets:

$$A_t = \{w_j, w_{j+1}, w_{j+2}\} \tag{4.2}$$

we calculate the cosine distance among the word embeddings as an average of the total number of instances of the token "farmer" across protest and counter-protest tweets to calculate the semantic polarity (SP):

$$SP = \frac{1 - CosineDistance(P_t, A_t)}{n_1 n_2} \tag{4.3}$$

where:

P_t : the set of word embeddings in the protest corpus

A_t : the set of word embeddings in the counter-protest corpus

w_i : the word embedding in the protest corpus

w_j : the word embedding in the counter-protest corpus

n_1 : the number of word embeddings in the protest corpus

n_2 : the number of word embeddings in the counter-protest corpus

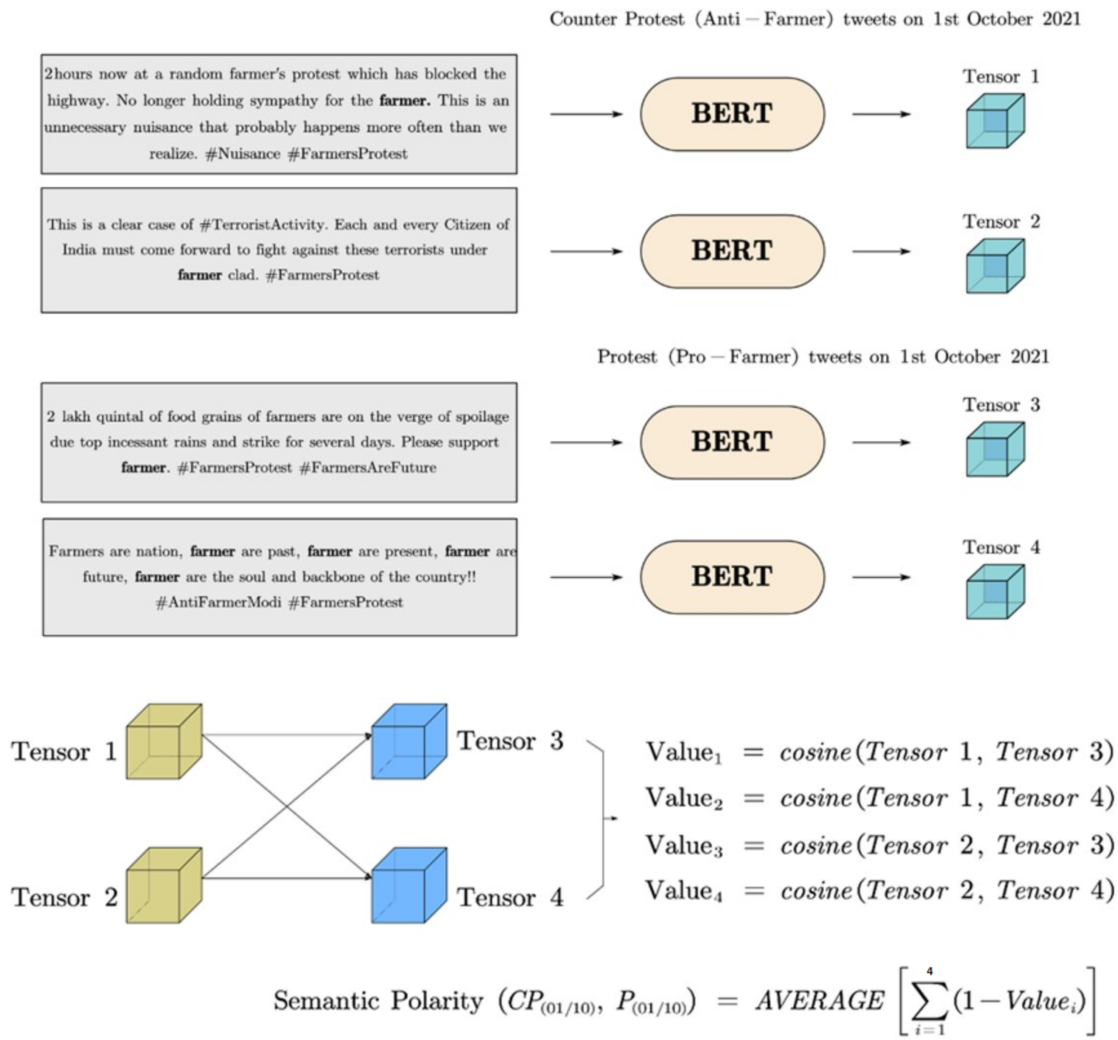

Figure 4.1: Calculating Semantic Polarization between Protest and Counter-Protest Tweets.

In Figure 4.1, we showcase the semantic polarization (SP) score in how the protest and counter-protest communities use the keyword, “farmer” during the Indian farmers’ protest. In the first step, we extract all contextual word embeddings (a fixed tensor of size 768) associated with every occurrence of the token “farmer” in our corpus on 1st October 2021. Next, we calculate the cosine distance between the tensors associated with the word representation "farmer" in the protest corpus and the counter-protest corpus. Finally, we take the average of all the cosine distance values to obtain the semantic polarization score or the semantic divergence in the usage of the token "farmer" between the protest and counter-protest communities on 1st October 2021.

4.5 RQ4: How Does the Use of Offspring Hashtags in Discussions Related to Farmers by Protest and Counter-Protest Communities Affect Semantic Polarization Trends?

Given our findings in RQ3, we used Granger causality [44] to test whether offspring hashtags significantly forecast semantic polarization in how the protest and counter-protest communities discuss farmers. Granger causality is a statistical method used to determine if one time series X is useful in predicting another time series Y (as described by Granger [44]). In simpler terms, if two time series X and Y are aligned, we can say that X Granger-causes Y if past values of X (denoted as X_{t-1} belongs to X) provide better predictions for the current value of Y (denoted as Y_t) than using only past values of Y (denoted as Y_{t-1} belongs to Y) alone [44]. The time point is represented by t, and the lag time is represented by l. Based on this idea, we test H1 and H2 mentioned in Chapter 1, Section 1.3 of this thesis.

Chapter 5

Results

5.1 Summary

Our study reveals that offspring hashtags from both the protest and counter-protest communities emerge almost immediately and rapidly, just mere minutes after the first appearance of the dominant and better-known parent hashtag (RQ1). In just 1 month into the protest, 13% of all the protest and 20% of all the counter-protest offspring hashtags from the entire protest period appeared. There are also significant differences in the framing and discussion of the keyword "farmer" between protest and counter-protest communities in posts that contain offspring hashtags (RQ2A). Protest communities use frames associated with recognition ("backbone", "heroes", "giver", "hard work"), reification ("food", "fields"), and empowerment ("unity", "win", "support") - essentially, frames that emphasize the social identity, occupation, and support of farmers. By contrast, the counter-protest group focuses on frames that delegitimize ("fake", "propaganda", "political tool") and politicize ("ISI", "BJP") farmers involved in the protest to bolster nationalistic views ("anti-national[1]", "Indian"). Both groups use offspring hashtags in a manner that reinforces frames that exacerbate rather than alleviate polarization between perspectives (RQ2B). We also calculate the semantic polarization or the semantic distance between how the two groups contextually discuss the keyword, "farmer" and find that it fluctuates with key protest events and changes in the volume of offspring hashtags (RQ3). Given these results, we test whether offspring hashtags significantly forecast semantic polarization by conducting two Granger-causality tests

[1]This token is used to refer to the farmers as miscreants working against the interests of the nation.

(RQ4). Our results show that the fluctuating volumes of offspring hashtags significantly predict semantic polarization, whereas the volume of parent hashtags does not. In other words, offspring hashtags play a more significant role in polarizing protest discourse between the protest and counter-protest communities.

5.2 Networked Growth of Offspring Hashtags (RQ1)

The networked growth of offspring hashtags, as observed in Figures 5.1 and 5.2, illustrates differences in how protest and counter-protest offspring hashtags evolved over time. We use the methodology discussed in Chapter 4, Section 4.1 to arrive at the results shown in this section. In the first minute, there is a relatively equal number of offspring hashtags for both communities, and even an increase in the counter-protest community. However, this trend dramatically shifts as time progresses. For instance, at the 12-hour mark, the protest community has over twice as many offspring hashtags as the counter-protest community (59 vs. 23). The disparity in growth becomes even more evident over time. At the one-week mark, the protest community's offspring hashtags have grown three times the amount of the counter-protest community's (896 vs. 306). This trend continues to amplify with time, with the protest community generating almost ten times the number of offspring hashtags at the one-month mark (4387 vs. 459). By the end of the protest period, the protest community cultivated over fourteen times the number of offspring hashtags than the counter-protest community (32664 vs. 2240) (see Table 5.1).

The consistent growth in the number of offspring hashtags in the protest community over time suggests a sustained and increasing engagement among its members. This could be due to a combination of factors such as the strength of their shared ideas, the effectiveness of their communication strategies, and the level of organization within the community. The growth pattern also implies that the protest community is continually evolving its narratives and strategies, as evidenced by the proliferation of new offspring hashtags. Conversely, the slower growth rate of offspring hashtags in the counter-protest community may indicate a lower level of engagement, less effective communication strategies, or less resonance of the counter-protest narratives within its community. The relative stagnation in the growth rate over time might suggest that the counter-protest community's messaging and strategies are not evolving or adapting as rapidly or as effectively as the protest community.

Time-Step	**Offspring Hashtags (Protest)**	**Offspring Hashtags (Counter-Protest)**
t_1 =1 min	1	8
t_2 = 1 hour	11	11
t_3 = 12 hrs.	59	23
t_4 = 24 hrs.	92	40
t_5 = 3 days	263	103
t_6 = 1 week	896	306
t_7 = 1 month	4,387	459
t_8 = 6 months	16,212	1,622
t_9 = 1 year	32,664	2,240

Table 5.1: Networked Growth of Offspring Hashtags Over Time.

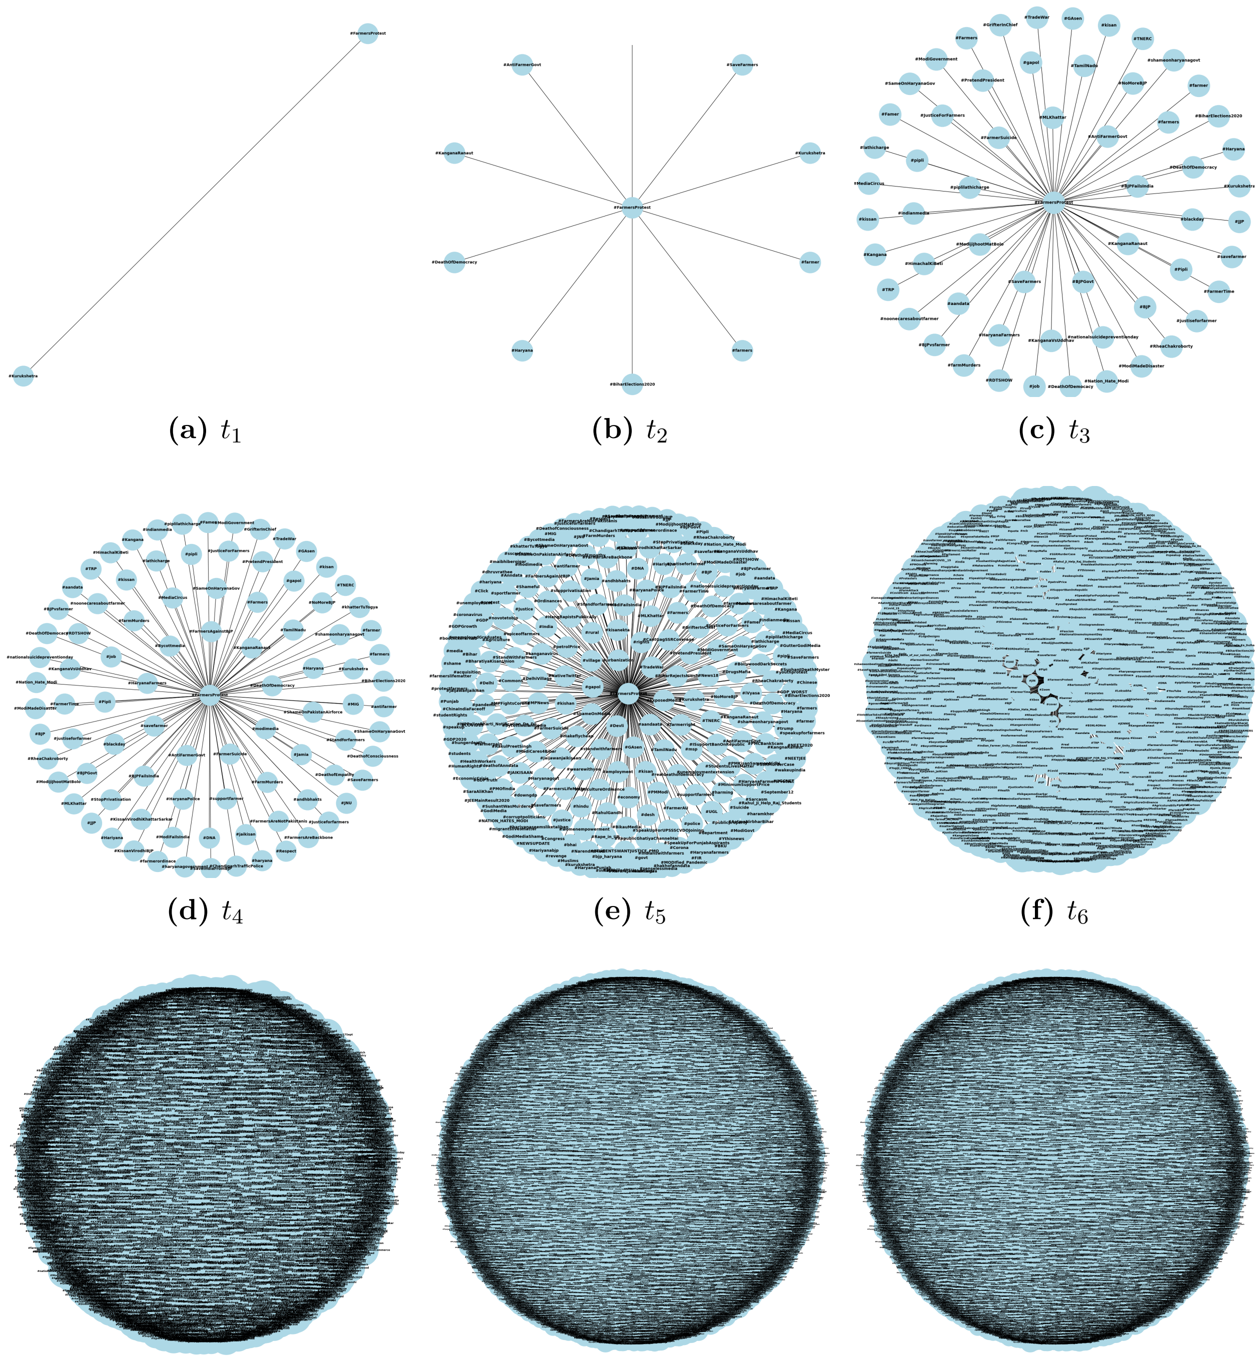

(a) t_1 **(b)** t_2 **(c)** t_3

(d) t_4 **(e)** t_5 **(f)** t_6

(g) t_7 **(h)** t_8 **(i)** t_9

Figure 5.1: Network Growth of Protest Offspring Hashtags.

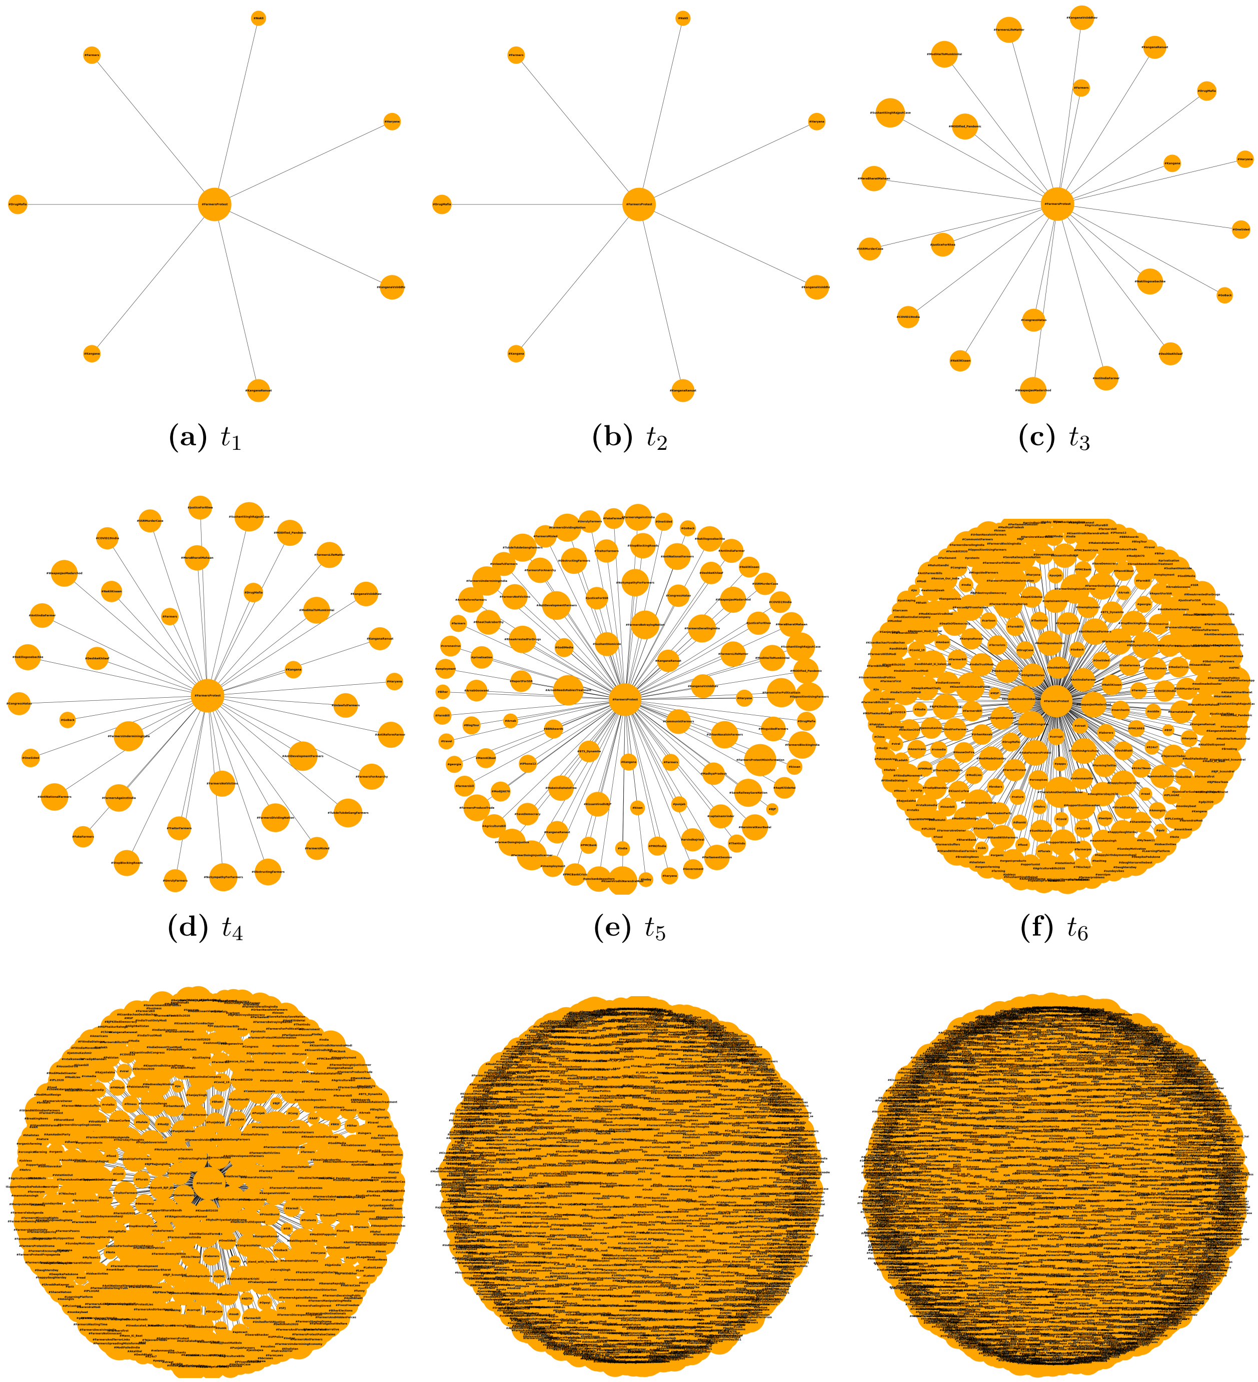

(a) t_1 (b) t_2 (c) t_3

(d) t_4 (e) t_5 (f) t_6

(g) t_7 (h) t_8 (i) t_9

Figure 5.2: Network Growth of Counter-Protest Offspring Hashtags.

5.3 Protests and Counter-Protests Significantly Diverge in Framing of Farmers (RQ2A), but Use Offspring Hashtags Similarly (RQ2B).

Protesters and counter-protesters who use offspring hashtags significantly differ in how they frame farmers in their posts. Table 5.2 shows the top 20 words that are contextually most similar to the token "farmer" based on descending order of cosine similarity values. These contextual tokens serve to evoke and highlight the various frames through which the two groups define farmers in the context of the protest. As shown, we categorized the tokens into relevant thematic frames most consistent with how the protest and counter-protest communities contextually used these words in their tweets to frame their perception of farmers. The frames put forth by the protesters and counter-protesters are in contention with each other. The protest communities use frames that recognize, reify, and empower farmers while counter-protesters delegitimize and politicize farmers to evoke nationalistic sentiment against those who support the farmers' protest. We delve more deeply into the results of our discourse analysis to understand how these contextual words are used to evoke such frames and how the two groups use offspring hashtags to solidify their framing of farmers. In our discourse analysis results (sub-section 5.3.1 to sub-section 5.3.3), we bold the contextual tokens and highlight the offspring hashtags in blue (protest) and red (counter-protest) in the tweets shown.

Counter-Protest Frames	Contextual Tokens	Protest Frames	Contextual Tokens
Delegitimization	Fake	Recognition	Backbone
	Khalistani		Heroes
	Hijacked		Giver
	Exposed		Distress
	Doublespeak		Earns
	Nihang Murderer[a]		Life
	Tilkait Exposed [b]		Outstanding
Politicization	Political Tool	Retification of Farmer Identity (Agriculture / Farming)	Fields
	Propaganda		Earns
	Farmers_WithPMModi		Foods
	Modi_With_Farmers		Hard Work
	BJP [c]		Distress
	ISI	Empowerment	Support
	Manmohanji [d]		Unity
Nationalism	Indian		**Zindabad** [e]
	Anti-National		Real
	DeshNahiBhulega [f]		Win
	Rihanna		Lives
	Justin Trudeau	Misc	Cute
	Mia Khalifa		Fun

Explanation of foreign keywords
a. "Nihang" refers to an armed Sikh warrior associated with anti-nationalist sentiment in India
b. Rakesh Tikait is a prominent farmer accused of corruption and extremism
c. BJP, which stands for Bharatiya Janata Party refers to a right-wing party in India
d. Former Prime Minister of the Indian National Congress (center-left party)
e. "Zindabad" is a Urdu word that means "long live" or "victory".
f. The phrase translates to "The country will never forget"

Table 5.2: Top 20 Words Most Contextually Similar to "farmer".

5.3.1 Recognition vs. Delegitimization

A prominent frame that emerged in the protest community's discussion of farmers was recognition, as demonstrated by their use of words such as "backbone", "heroes", and "giver", as noted in Table 5.2. These words convey a sense of appreciation and respect for the crucial role that farmers play in society. In their tweet, one pro-farmer poster emphasized the importance of farming in supporting other economic sectors, asserting that assisting farmers is tantamount to safeguarding the country and humanity at large.

*"Farming is the **backbone** of other industries. Save farmer, save humanity, save country. #RepealFarmLawsNOW #FarmersProtest"*

Here, the poster's use of the word "backbone" highlights the farmers' role in keeping the Indian economy afloat. The poster also employs the offspring hashtag #RepealFarmLawsNOW, which serves to urgently draw attention to the struggles faced by farmers due to the recent agricultural laws. The repeated use of the word "save" further highlights the earnestness of the situation and the crucial role of farmers in ensuring the country's prosperity. Another user echoes similar sentiment towards the vitality of Indian farmers in their tweet:

*"Indian farmers are one of the most important members of the society as they are the producer and **giver** of the food for whole country. #NoFarmersNoFuture #FarmersProtest"*

The overall tone of the statement is one of respect and appreciation for farmers, recognizing their indispensable role as "givers" to society. The protester uses the offspring hashtag #NoFarmersNoFuture as a form of meta-commentary to reinforce the importance of farmers, suggesting that their welfare is critical to the future of the country. By contrast, counter-protesters seek to delegitimize farmers or protest participants in their discussion by denouncing them as fake farmers or associating them with rapists and terrorist organizations. For example:

*"R@pe at Tikri Border??? Well, thank a **fake** farmer!!! #FarmersProtest #BalatkariAndolan"*

Here, the poster refers to an alleged rape incident of a 25-year old woman who participated in the anti-farm law rally at a protest site near the Tikri border in Delhi [88]. While the incident could not be factually verified, the counter-protester insinuates that the alleged crime can be attributed to those pretending to be farmers ("thank a fake farmer!!!"). The poster takes a step further, implying that protest participants, including the farmers, are rapists: the offspring hashtag, #BalatkariAndola translates to protest of rapists ("Balatkari" means rapist and "Andolan" means protest in Hindi).

"Indira thhok Di, Modi bhi Thhok denge Such is the level of Congress and their sponsored Farmers. Shame on all Wokes, Commies and #UrbanNaxals who support these Khalistani ***Fake*** *Farmers. #FarmersProtest"*

Here, another counter-protester uses strong, provocative language to discredit the protesting farmers and their supporters. The phrase "Indira thhok Di, Modi bhi Thhok denge" (Translation: Killed Indira Gandhi, We will kill Modi too) refers to the assassination of former Indian Prime Minister Indira Gandhi and insinuates that the "sponsored Farmers" are willing to resort to violence against the current Prime Minister, Narendra Modi. "Urban Naxals" is a derogatory slang commonly used in the context of Indian politics when referring to someone with alleged ties with communist militants or left-wing socialists that are hostile to Hindu nationalism. The counter-protester uses the offspring hashtag, #UrbanNaxals, along with expressions, such as "Wokes" and "Commies" to frame farmers as illegitimate, "Khalistani" secessionists: Khalistan is a political ideology that led to a violent separatist movement led by the Sikhs in the 1980s. As such, a large number of counter-protesters in our data framed farmers as Khalistanis or Urban Naxals who were participating in the protest to destabilize the Indian government.

5.3.2 Reification vs. Politicization

Another frame that emerged in the protest community's discussion of farmers was reification of farmer identity, as demonstrated by their use of words such as "hard work", "fields" and "foods" as noted in Table 5.2. These words underscore the importance of farmers' contributions to society, emphasizing their dedication, connection to the land, and role in food production.

"No matter how digital India becomes, roti will not be downloaded from Google.. That will be achieved only by the ***hard work*** *of the farmer. I support farmers #TanashahiModiKisaanVirodhi #FarmersProtest"*

The tweet highlights the significance of farmers' "hard work" in food production and how technology cannot replace manual agricultural labor. However, the use of the offspring hashtag #TanashahiModiKisaanVirodhi (Translation: Dictator Modi Anti Farmers), detracts from the tweet's main message and directs attention towards a potentially divisive political debate, weakening the core message of the farmers' protest.

Alternatively, those opposing the farmers' protest sought to politicize the farmers' protest by linking them to external forces and alleged misinformation. They utilized language that included terms like "propaganda" and "ISI" (referring to Inter-Services Intelligence, Pakistan's foreign intelligence agency), suggesting ulterior motives behind the protest and implying that the movement aimed to undermine India's stability.

"Pakistan's intelligence agency ***ISI*** *has hijacked the farmer protest. It has been proved in all standards of the scales. #ISI #FarmersProtest"*

The counter-protester politicizes the farmers' protest by referring to "ISI" to imply the potential influence by a foreign intelligence agency from a country that has historically shared diplomatic tensions with India. The user also takes an absolutist tone, through expressions, such as "It has been proved in all standards of scales" to solidify their stance, leaving little room for alternative perspectives. The offspring hashtag #ISI reinforces the tweet's core message, possibly influencing public opinion by connecting the protests to external threats.

5.3.3 Empowerment vs. Nationalism

The third frame that emerged from our analysis of the protest tweets was a call to empower farmers. The protest community emphasized the need to stand in solidarity with farmers, by using contextual words that evoke support to empowering farmers, such as "unity", "win", and "support". As shown in the tweet below, the protester hints a hopeful encouragement that worldwide support for the farmers and the protest is internationally picking up, defining such solidarity as "unity":

"The support internationally is picking up, this is ***unity****, whilst some of India's celebrities are Modi puppets. #HitlerModi Against Farmers Boycott BJP Support Farmers #FarmersProtest"*

This positive and hopeful tone is drastically contrasted with the offspring hashtag "#HitlerModi". In fact, throughout our analysis of the protest tweets, we found protesters often framed the protest as a fight between farmers fighting for their lives under an oppressive government similar to that of Hitler:

"Struggle of farmers WON! Ego of #HitlerLost"

As evidenced by the words linked with optimism, solidarity and support (e.g., "support", "unity", "win", "won"), protesters tend to highlight the empowerment of farmers in their framing of the protest and the farmers. However, in doing so, they also tend to discredit and invalidate those who oppose the protest or support the government's laws in a manner that diverts attention from the core issues of the protest, sometimes even resorting to extreme comparisons such as equating them with Hitler as demonstrated through offspring hashtags, such as #HitlerModi and #HitlerLost. The counter-protesters evoke nationalistic frames by using offspring hashtags in a similar manner that shifts focus away from the issues of the protest. For example:

"This is what #FarmersProtest did to India by the #Canada based separatist #khalistani agents ***Justin Trudeau*** *came out in support of the divisive #Khalistan agenda AGAINST #indian government #CriticalRaceTheory - when laws ideals apply differently to different races!"*

In this tweet, the counter-protester uses the offspring hashtag #CriticalRaceTheory to evoke a controversial fixture in the American debate around how schools should grapple with the history of race and racism in the United States. While an important issue, the offspring hashtag, #CriticalRaceTheory is off-topic from the farmers' protest, turning the protest into an issue of race: "when laws apply differently to different races!". The counter-protester uses the offspring hashtag #CriticalRaceTheory to depict the protest as an instance of unequal treatment or double standards, contributing to divisions along regional, religious, and political lines, and thereby destabilizing the nation. Similarly, the user mentions the Canadian Prime Minister Justin Trudeau (a top contextual token in our counter-protest data), which serves a similar purpose in framing the farmers' protest as a threat to nationalistic values.

Overall, both the protest and the counter-protest communities use offspring hashtags in a manner that reinforces frames that exacerbate rather than alleviate polarization between perspectives. Both sides use offspring hashtags to evoke conspiracist notions, vehemently malign opponents (e.g., rapists, Hitler), and invalidate any counterarguments in an uncompromising manner. As a result, this leaves little room for bridging perspectives or the potential for substantive discussions on the actual issues of the protest that are critical to the livelihood of the farmers. Unfortunately, offspring hashtags contribute to reducing the protest dialogue to sensationalized soundbites and misleading fallacies.

5.4 Semantic Polarization in How the Protests vs. Counter-Protests Discuss the Word “farmer” Evolves Closely With the Volume of Offspring Hashtags (RQ3)

To capture how offspring hashtags contribute to the linguistic polarization between the protest and counter-protest discourse at scale, we calculated the semantic polarization between the two groups’ discussion of the keyword “farmer” across posts with offspring hashtags evolves over time. Our semantic polarization calculation shows how discourse between protest and counter-protest communities about farmers closely evolves with the volume of offspring hashtags. Figure 5.3 shows how semantic polarization between the two groups that use offspring hashtags evolved from September 2020 to November 2021. Semantic polarization ranged from 0.59 to 0.82 (values closer to one indicate high polarity and values closer to zero indicate low polarity [31]). Semantic polarity started off high at the beginning of the protest in the fall of 2020, reaching a dip in May 2021, and then rapidly ascended after July 2021 (orange line, Figure 5.3). This pattern evolves closely with the fluctuating volume of offspring hashtags. To compare whether this pattern holds true in how protesters and counter-protesters that do not use offspring hashtags in their posts, we replicated the analysis using tweets that only contained parent and no offspring hashtags **(see Appendix B for details)**. As shown in **Figure B.1 (Appendix B)**, the semantic polarity in how protesters vs. counter-protesters who only use the parent tag, #FarmersProtest and no offspring hashtags, does not fluctuate as drastically.

Furthermore, the daily average semantic polarity between the two opposing groups that do not use offspring hashtags is 0.54 whereas the daily average semantic polarity between the protesters vs. counter-protesters that use offspring hashtags is 0.74 - approximately 37.4 % higher on average in **Table B.3 (Appendix B)**. Unlike the close parallel pattern between the semantic polarity trends and the fluctuating volume of offspring hashtags throughout the protest period across those who use offspring hashtags, semantic polarization trends across those who only use parent tags do not seem to follow the volume changes in tweets that only contain parent tags in **Figure B.1 (Appendix B)**. The findings of RQ3 not only demonstrate how semantic polarization evolves across protest and counter-protest communities that use offspring hashtags, but also that this pattern evolves closely with the volume of offspring hashtags that emerge over time as well. Yet, this case only holds true for offspring and not for parent hashtags, implying that the impact and role of offspring hashtags in polarizing online public discourse during a social movement may be much more significant than that of just the parent hashtags alone.

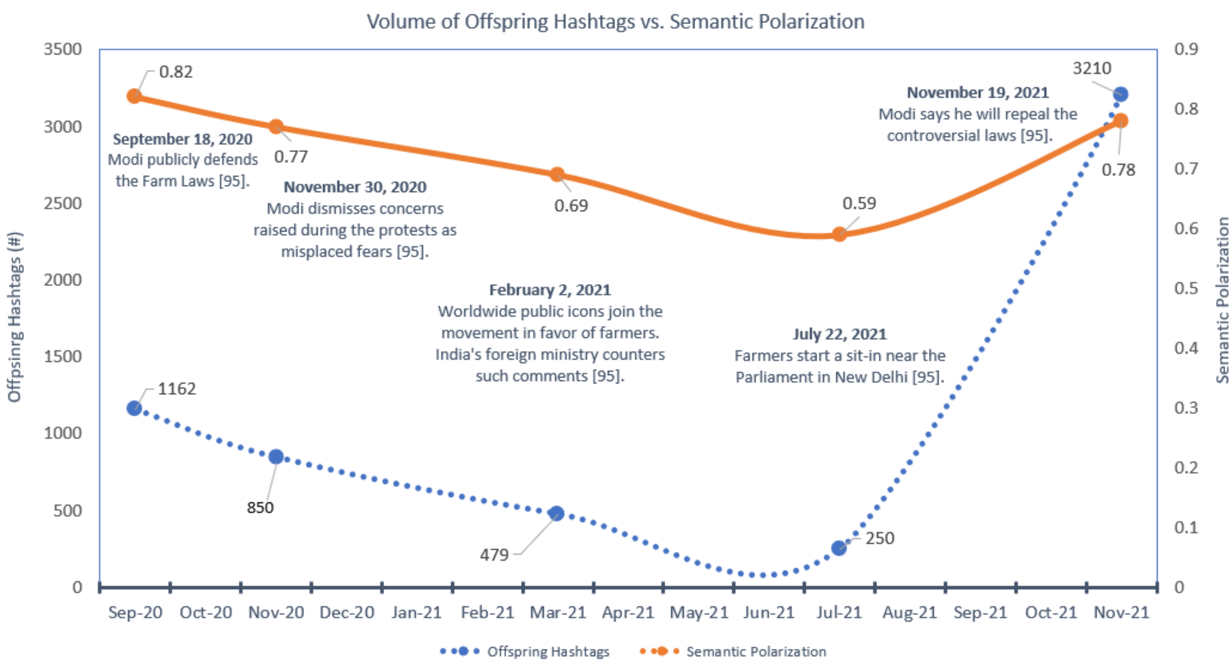

Figure 5.3: Evolution of Semantic Polarity Around the Word "farmer" Between Protests and Counter-Protests That Use Offspring Hashtags.

5.5 Trends in Fluctuating Volumes of Offspring Hashtags Significantly Forecast Semantic Polarization Patterns in How the Protest and Counter-Protest Communities Discuss Farmers (RQ4).

Given our results in RQ3, we use Granger-causality to test whether the fluctuating volumes of offspring hashtags throughout the protest is in fact a significant predictor of how semantic polarization between the protest vs. counter-protest communities evolves over time. In this section, we present the outcomes of the Granger causality results. We observe that the volume of offspring hashtags over time can significantly forecast the shift in semantic polarity in how the two groups discuss the keyword “farmer”. Granger causality results for Hypotheses 1 and 2 with corresponding significant lag lengths ($p < 0.05$) are shown in Table 5.3. To ensure the results obtained through Granger-causality tests were not merely a function of time, we also conducted the Augmented Dickey-Fuller (ADF) test [21]. We checked for stationarity of time series using the ADF test (**see Table A.1 and Table A.2 (Appendix A)** before running our Granger causality analysis.

H1*: The daily number of offspring hashtags significantly Granger-causes semantic polarization between the protest and counter-protest communities' discussion of farmers (lag= 2 days; p-value = 0.0473).*

H1. As highlighted in Table 5.3, the daily count of offspring hashtags used by protest and counter-protest communities significantly forecasts the shift in semantic polarity around the token “farmer” with a time lag of 2 days. This suggests that it takes 2 days for the offspring hashtags used by protests and counter-protests to forecast semantic polarity around the token “farmer”.

H2: *The daily number of offspring hashtags that newly emerge each day significantly Granger-causes semantic polarization between the protests and counter-protest communities' discussion of farmers (lag= 3 days; p-value = 0.0431).*

H2. As highlighted in Table 5.3, the daily number of offspring hashtags that newly emerge each day significantly forecasts semantic polarization in protest and counter-protest communities online. Our results show that it takes merely 3 days for newly emerged offspring hashtags to influence discourse around how the token “farmer” diverges between protest and counter-protest communities online.

Testing for Unidirectional Causality. To test whether the direction of causality was such that the volume of hashtags influenced semantic polarization (offspring hashtag volume $\rightarrow$ semantic polarization), and not vice-versa, we tested the reversed hypotheses for H1 and H2 (semantic polarization $\rightarrow$ offspring hashtag volume) and found the **reverse of H1 and H2 to be non-significant** (H1 reverse: p-value = 0.865; H2 reverse: p-value = 0.741). In other words, the volume of offspring hashtags impact how the protest and counter-protest communities diverge in their discussion of farmers over the course of the protest, and not vice-versa.

Testing for Causality in Tweets that Only Contain Parent Hashtags. Furthermore, to test whether the volume of tweets that only contained the parent hashtag and not offspring hashtags contributed to the patterns in which the protests and counter-protests semantically diverged in their discussion of farmers, we tested Hypothesis 3 (H3):

H3: *The volume of tweets that exclusively contain the parent hashtag significantly Granger-causes semantic polarization between protests and counter-protest communities' discussion of farmers.*

H3: As highlighted in Table 5.3, we observe that the Granger-causality relationship is **non-significant**. In other words, the volume of tweets that exclusively contain the parent hashtag **do not** affect how the protest and counter-protest communities diverge in their discussion of farmers.

Hypothesis	Lag value (days)	P-value
H1. The daily unique count of offspring hashtags significantly Granger-causes semantic polarization between protests and counter-protest communities' discussion of farmers.	2	0.0473
H2. The unique count of offspring hashtags that newly emerge each day significantly Granger-causes semantic polarization between protests and counter-protest communities' discussion of farmers.	3	0.0431
H3. The volume of tweets that exclusively contain the parent hashtag significantly Granger-causes semantic polarization between protests and counter-protest communities' discussion of farmers.	3	0.982

Table 5.3: H1, H2 and H3 Results Using Granger Causality Tests.

Chapter 6

Discussion

6.1 Framing Effects on Marginalized Groups in Online Social Movements

Understanding how key social groups are discussed in online social movements (e.g., “police”, “blacks” in #BlackLivesMatter, “women” in #MeToo, etc.) can shed critical light on how these groups are impacted by social media protests [16, 27]. For example, public attitudes towards law enforcement and black individuals have radically shifted and polarized through the waves of #BlackLivesMatter movements over the past several years [14]. Similarly, the #MeToo movement has also induced changes in people’s attitudes towards women and their experiences of sexual harassment [45]. Social groups at the center of protests are typically marginalized or belong to minoritized groups [112, 120] and are heavily impacted by protest dynamics and outcomes [98, 115]. Conversations and events surrounding protests also strongly influence the public’s attitudes towards these minority groups [17, 63, 99]. Hence, in this research, we use the Indian farmers’ protest as a contextual lens to understand how language and framing patterns evolve through the usage of offspring hashtags and how they influence discourse online [19].

Our embedding analysis in RQ2 revealed significant differences in the most common tokens associated with the term “farmer” between protests and counter-protests. These discrepancies suggest potential framing effects [34, 70] in how the #FarmersProtest is discussed across online communities supporting protests and counter-protests. Frames in social movements are often generated through distributed framing [65], where any user can contribute to shaping the dominant narrative. As theorist Jacques Ellul observed, propaganda aims to steadily reinforce or undermine stereotypes, working best when it connects to the fundamental currents of the society it seeks to influence [33].

Both the protests and counter-protests seek to tap into powerful social currents [10] to frame narratives around the Indian farmers' protest. Indian society has been particularly responsive to issues of social justice, especially when injustice is framed as a threat to Indian culture and identity [12]. Moreover, the government's introduction of new farm laws was seen as a betrayal of farmers, leading to widespread anger and resentment [48]. The protest community capitalized on this sentiment by framing the protest as an attack on human rights and highlighting the sufferings of the farmers, thereby mobilizing additional support for their cause [53]. Conversely, the counter-protest framing appeals to long-established fears of the "enemy within" in India, such as the fear of foreign influence and terrorism [41, 55]. The idea that Indian society is being controlled and manipulated by foreign governments has long played a role in conservative theorists' narrative [55].

This stark contrast in the linguistic tone and style of protests and counter-protests, as they develop contesting frames around the Indian farmers' protest, not only leads to semantic divergence between the two communities in their understanding of the movement but also results in users developing alternate views of reality surrounding the same events [53, 54, 79].

6.2 Understanding Polarization in Protest Discourse: Offspring Hashtags and Moral Panic

As online social movements gain momentum, a large number of people create offspring hashtags to share their opinions and views, which often results in the fragmentation and dilution of the core message [71]. This was highlighted from our analysis in RQ4, where the volume and uniqueness of offspring hashtags were shown to significantly affect the shift in semantic polarization of discourse between protests and counter-protests. In such cases, offspring hashtags often present skewed and exaggerated views of reality that don't align with actual events (Table 1.1). Such offspring hashtags gain popularity and drive discourse around the farmers' protest, overshadowing the original motivation of the protest [81].

Our observation emphasizes that the increased usage and uniqueness of offspring hashtags, coupled with their role in polarizing narratives, positions them as identity-based socio-political movements rooted in moral panic [22] rather than facilitating constructive discourse between divergent discourse communities [81]. Moral panic is a theory in sociology developed by Cohen [22] that describes a widespread fear, most often an irrational one, that someone or something is a threat to the values, safety, and interests of a community or society at large [22]. Cohen highlights the cause of moral panic as a "condition, episode, person or group of persons emerges to become defined as a threat to societal values and interests" [22]. Moral panic arises when distorted communication campaigns are used to create fear and reinforce stereotypes based on race, ethnicity, and social class [119]. Factors such as deviation from core issues (#ModiKisaanVirodhi vs. #RahulPriyankaFakeDrama) [71], emotional manipulation (#FarmerLivesMatter vs. #AntiNationalFarmers) [81], formation of echo chambers (#NationWithFarmers vs. #FarmersAgainstNation), group identity (#IStandWithFarmers vs. #FakeFarmers) [22], and amplification of misinformation and conspiracy theories (#ModiHataoDeshbachao vs. #AtyachaariSarkaar) [22] all contribute to the characterization of discourse embedded in moral panic. These factors highlight that networked moral panics, fueled by offspring hashtags, can distort information and exaggerate reality to disseminate misinformation under the guise of activism.

As a result, promoting balanced information, critical thinking, and media literacy in online spaces is essential. In addition, encouraging constructive dialogue and developing fact-checking algorithms can help mitigate the consequences of distorted communication and cultivate a more balanced understanding of critical issues online.

6.3 The Future of Offspring Hashtags in India: A Double-Edged Sword for Online Social Movements

India's history has been significantly shaped by citizen-led protests, which have played a crucial role in the country's democratization [41]. However, recent developments have given rise to concerns about the efficacy of public participation in social issues, particularly in the context of online discourse. In this light, we explore the double-edged nature of offspring hashtags, which can both empower and endanger social movements. India's social media landscape has been marked by a rising prevalence of online misinformation and polarization [39]. The government's increasing influence over both the press and social media platforms [102], coupled with attempts to manipulate information [24, 25, 48], has exacerbated challenges faced by social movements [64, 85, 86]. Furthermore, low literacy rates and unrestricted access to social media have contributed to polarizing the uneducated youth and exacerbating discrimination against marginalized groups [11, 41].

On the positive side, offspring hashtags can amplify voices, garner international attention, and mobilize support for social movements [10, 53]. The success of the farmers' protest was largely attributed to the role of hashtags and social media platforms [10, 53, 54]. However, the negative side of offspring hashtags is equally potent. They can polarize public opinion, incite violence, and spread misinformation. A case in point is the Citizenship Amendment Act (CAA) protests, where political hashtags propagated hostility and suspicion toward protesting social groups [85]. This online hostility spilled over into the real world, resulting in widespread violence and riots in February 2020 [85].

The future of offspring hashtags in India thus presents a delicate balance of opportunities and challenges. While they can empower grassroots campaigns [53, 54] and amplify marginalized voices [66], there is a risk of them being weaponized for polarizing public opinion, inciting violence, and spreading misinformation [85]. Therefore, it is crucial for stakeholders, including social media platforms, civil society organizations, and individuals, to actively monitor the use of offspring hashtags. This includes fostering responsible online behavior by promoting a culture of fact-checking [29] and critical thinking, enhancing digital literacy through education and awareness programs [8], and implementing effective content moderation policies [94] that curb the spread of misinformation [29]. By doing so, India can uphold the democratic nature of online discussions and ensure that social movements continue to drive meaningful progress and reform.

Chapter 7

Limitations

Our research focuses on the analysis of social media discussions surrounding the Indian farmers' protest, and the results may not be universally applicable to other online social movements. Additionally, our methodology may face limitations when dealing with multilingual data [92] or code-switching [106], which could potentially impact the assessment of semantic polarization in these conversations. Lastly, although the imbalanced nature of our dataset may not significantly affect the determination of themes, since the methods used for calculating findings for RQ2 focus on the relative frequencies and distribution of offspring hashtags within each group rather than directly comparing the absolute number of tweets between protest and counter-protest, it still presents a limitation in accurately assessing the extent of semantic polarization between protests and counter-protests using offspring hashtags. This could result in the counter-protest community's embeddings failing to fully capture the diverse meanings and nuances associated with the term "farmer". Consequently, for future work, we would explore strategies for addressing class imbalance, such as re-sampling techniques [18] or weighting schemes [84], to enhance the reliability of semantic polarization measurements and ensure a more comprehensive understanding of the discourse dynamics between protest and counter-protest communities.

Chapter 8

Conclusions

In this thesis, we have investigated the role of offspring hashtags on the semantic polarization of online discourse between the protest and counter-protest communities during the 2021 farmers' protest in India. Our findings reveal that the growth of protest community offspring hashtags was significantly faster and more expansive than that of the counter-protest community, implying a higher engagement and evolving narratives within the protest community (RQ1). Moreover, both protests and counter-protests use offspring hashtags in a manner that further polarizes rather than bridges perspectives on core issues, focusing on themes that malign the other side (RQ2). We also found that the semantic polarization in how these two groups talk about farmers evolves closely with key protest events as well as the volume of offspring hashtags (RQ3). In fact, the fluctuating volume of offspring hashtags throughout the protest period is a significant and reliable predictor of semantic polarization patterns in how the protests and counter-protests diverge in their framing and discussion of farmers (RQ4). Our findings have significant implications for understanding the impact of offspring hashtags created by protest and counter-protest groups on polarizing networked discussions between the two communities. Our research sheds light on how online protest hashtags can potentially contribute to the spread of extremist views, further exacerbating the division between opposing protest camps rather than bridging it. Future research should explore the influence of images and videos accompanying offspring hashtags on polarization and identify which offspring hashtags serve as tipping points that result in heightened discourse polarization.

Bibliography

[1] Lahbib Ajallouda, Kawtar Najmani, Ahmed Zellou, and El habib Benlahmar. Doc2Vec, SBERT, InferSent, and USE Which embedding technique for noun phrases? In *2022 2nd International Conference on Innovative Research in Applied Science, Engineering and Technology (IRASET)*. IEEE, 2022.

[2] Sohail Akhtar, Valerio Basile, and Viviana Patti. Modeling Annotator Perspective and Polarized Opinions to Improve Hate Speech Detection. In *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, volume 8, pages 151–154, 2020.

[3] Neha Alawadhi. Twitter trends: 2021's good, bad, happy and sad through hashtags. *Business Standard*, 2021. URL https://www.business-standard.com/article/current-affairs/twitter-trends-2021-s-good-bad-happy-and-sad-through-hashtags-121120900076_1.html. Accessed January 12, 2023.

[4] Nuha Albadi, Maram Kurdi, and Shivakant Mishra. Hateful People or Hateful Bots?: Detection and Characterization of Bots Spreading Religious Hatred in Arabic Social Media. *Proceedings of the ACM on Human-Computer Interaction*, 3(CSCW):1–25, 2019.

[5] Rabab Alkhalifa and Arkaitz Zubiaga. Capturing stance dynamics in social media: open challenges and research directions. *International Journal of Digital Humanities*, 3(1-3):115–135, 2022.

[6] Nazanin Andalibi, Ashley Lacombe-Duncan, Lee Roosevelt, Kylie Wojciechowski, and Cameron Giniel. LGBTQ Persons' Use of Online Spaces to Navigate Conception,

Pregnancy, and Pregnancy Loss: An Intersectional Approach. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 29(1):1–46, 2022.

[7] Ahmer Arif, Leo Graiden Stewart, and Kate Starbird. Acting the Part: Examining Information Operations Within #BlackLivesMatter Discourse. *Proceedings of the ACM on Human-Computer Interaction*, 2(CSCW):1–27, 2018.

[8] R Babu, S Kalaivani, and K Saileela. *Empowering India Through Digital Literacy (Vol. 1)*. Lulu.com, 2019.

[9] Duilio Balsamo, Valeria Gelardi, Chengyuan Han, Daniele Rama, Abhishek Samantray, Claudia Zucca, and Michele Starnini. Inside the Echo Chamber: Disentangling network dynamics from polarization. *arXiv preprint arXiv:1906.09076*, 2019.

[10] Orin Basu. Battle of the hashtags: Debate over farmers' protest goes global for local. *India Today*, February 3 2021. URL https://www.indiatoday.in/news-analysis/story/battle-of-the-hashtags-debate-over-farmers-protest-goes-global-for-local-1765657-2021-02-03. Accessed May 9, 2023.

[11] Natasha Behl. India's farmers' protest: an inclusive vision of Indian democracy. *American Political Science Review*, 116(3):1141–1146, 2022.

[12] Sunil Bhatia. *Decolonizing Psychology: Globalization, Social Justice, and Indian Youth Identities*. Oxford University Press, 2017.

[13] Anand Bihari and Manoj Kumar Pandia. Eigenvector centrality and its application in research professionals' relationship network. In *2015 International Conference on Futuristic Trends on Computational Analysis and Knowledge Management (ABLAZE)*, pages 510–514. IEEE, 2015.

[14] Jeffrey Layne Blevins, James Jaehoon Lee, Erin E McCabe, and Ezra Edgerton. Tweet-

ing for social justice in #Ferguson: Affective discourse in Twitter hashtags. *New Media & Society*, 21(7):1636–1653, 2019.

[15] Kyle Booten. Hashtag Drift: Tracing the Evolving Uses of Political Hashtags Over Time. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, pages 2401–2405, 2016.

[16] Maia J Boyd, Jamar L Sullivan Jr, Marshini Chetty, and Blase Ur. Understanding the Security and Privacy Advice Given to Black Lives Matter Protesters. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, pages 1–18, 2021.

[17] Laura Broasca, Versavia-Maria Ancusa, and Horia Ciocarlie. A Qualitative Analysis on Force Directed Network Visualization Tools in the Context of Large Complex Networks. In *2019 23rd International Conference on System Theory, Control and Computing (ICSTCC)*, pages 656–661. IEEE, 2019.

[18] Gregorius Satia Budhi, Raymond Chiong, and Zuli Wang. Resampling imbalanced data to detect fake reviews using machine learning classifiers and textual-based features. *Multimedia Tools and Applications*, 80:13079–13097, 2021.

[19] Dallas Card, Justin H Gross, Amber Boydstun, and Noah A Smith. Analyzing Framing through the Casts of Characters in the News. In *Proceedings of the 2016 conference on empirical methods in natural language processing*, pages 1410–1420, 2016.

[20] Nuno Carrico and Paulo Quaresma. Sentence Embeddings and Sentence Similarity for Portuguese FAQs. In *Proceedings of IberSPEECH 2021*, pages 200–204, 03 2021.

[21] Yin-Wong Cheung and Kon S Lai. Lag order and critical values of the augmented Dickey–Fuller test. *Journal of Business & Economic Statistics*, 13(3):277–280, 1995.

[22] Stanley Cohen. *Folk devils and Moral Panics*. Routledge, 2011.

[23] Joana Costa, Catarina Silva, Mário Antunes, and Bernardete Ribeiro. Concept Drift Awareness in Twitter Streams. In *2014 13th International Conference on Machine Learning and Applications*, pages 294–299. IEEE, 2014.

[24] Saloni Dash, Gazal Shekhawat, Syeda Zainab Akbar, and Joyojeet Pal. Extremism & Whataboutism: A Case Study on Bangalore Riots. *arXiv preprint arXiv:2109.10526*, 2021.

[25] Saloni Dash, Rynaa Grover, Gazal Shekhawat, Sukhnidh Kaur, Dibyendu Mishra, and Joyojeet Pal. Insights Into Incitement: A Computational Perspective on Dangerous Speech on Twitter in India. In *ACM SIGCAS/SIGCHI Conference on Computing and Sustainable Societies (COMPASS)*, pages 103–121, 2022.

[26] Saloni Dash, Dibyendu Mishra, Gazal Shekhawat, and Joyojeet Pal. Divided We Rule: Influencer Polarization on Twitter During Political Crises in India. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 16, pages 135–146, 2022.

[27] Munmun De Choudhury, Shagun Jhaver, Benjamin Sugar, and Ingmar Weber. Social Media Participation in an Activist Movement for Racial Equality. In *Tenth International AAAI Conference on Web and Social Media*, 2016.

[28] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. *arXiv preprint arXiv:1810.04805*, 2018.

[29] Apoorva Dhawan, Malvika Bhalla, Deeksha Arora, Rishabh Kaushal, and Ponnurangam Kumaraguru. FakeNewsIndia: A benchmark dataset of fake news incidents in India, collection methodology and impact assessment in social media. *Computer Communications*, 185:130–141, 2022.

[30] Mandeep Rai Dhillon. Opinion: The world's largest protests you've probably never heard of. *CNN*, 2021. URL https://www.cnn.com/2021/03/06/opinions/india-farmer-protests-mandeep-rai-dhillon/index.html. Accessed December 12, 2022.

[31] Xiaohan Ding, Mike Horning, and Eugenia H. Rho. Same Words, Different Meanings: Semantic Polarization in Broadcast Media Language Forecasts Polarity in Online Public Discourse. In *Proceedings of the International AAAI Conference on Web and Social Media*, 2023.

[32] Brianna Dym, Jed R Brubaker, Casey Fiesler, and Bryan Semaan. "Coming Out Okay" Community Narratives for LGBTQ Identity Recovery Work. *Proceedings of the ACM on Human-Computer Interaction*, 3(CSCW):1–28, 2019.

[33] Jacques Ellul, RK Merton, and Jean Lerner. *Propaganda: the formation of men's attitudes*. Vintage books, 1973.

[34] Robert M Entman and Nikki Usher. Framing in a Fractured Democracy: Impacts of Digital Technology on Ideology, Power and Cascading Network Activation. *Journal of Communication*, 68(2):298–308, 2018.

[35] Sheena Erete, Yolanda A Rankin, and Jakita O Thomas. I Can't Breathe: Reflections from Black Women in CSCW and HCI. *Proceedings of the ACM on Human-Computer Interaction*, 4(CSCW3):1–23, 2021.

[36] Ryan J Gallagher, Andrew J Reagan, Christopher M Danforth, and Peter Sheridan Dodds. Divergent discourse between protests and counter-protests: #BlackLivesMatter and #AllLivesMatter. *PloS one*, 13(4):e0195644, 2018.

[37] Ryan J Gallagher, Elizabeth Stowell, Andrea G Parker, and Brooke Foucault Welles. Reclaiming Stigmatized Narratives: The Networked Disclosure Landscape of #MeToo. *Proceedings of the ACM on Human-Computer Interaction*, 3(CSCW):1–30, 2019.

[38] Yun Gao, Hirokazu Hasegawa, Yukiko Yamaguchi, and Hajime Shimada. Malware Detection using Attributed CFG Generated by Pre-trained Language Model with Graph Isomorphism Network. In *2022 IEEE 46th Annual Computers, Software, and Applications Conference (COMPSAC)*, pages 1495–1501. IEEE, 2022.

[39] Kiran Garimella, Tim Smith, Rebecca Weiss, and Robert West. Political Polarization in Online News Consumption. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 15, pages 152–162, 2021.

[40] James Paul Gee. *An Introduction to Discourse Analysis: Theory and Method*. Routledge, 2004.

[41] Vindu Goel, Jeffrey Gettleman, and Saumya Khandelwal. Under Modi, India's Press Is Not So Free Anymore. *The New York Times*, 2, 2020. URL https://www.nytimes.com/2020/04/02/world/asia/modi-india-press-media.html. Accessed January 25, 2022.

[42] Andrea Gorra and Marianna Kornilaki. Grounded theory: experiences of two studies with a focus on axial coding and the use of the NVivo qualitative analysis software. *Methodology: Innovative approaches to research*, 1:30–32, 2010.

[43] Mark Gradstein and Maurice Schiff. The political economy of social exclusion, with implications for immigration policy. *Journal of Population Economics*, 19:327–344, 2006.

[44] Clive WJ Granger. Some recent development in a concept of causality. *Journal of econometrics*, 39(1-2):199–211, 1988.

[45] Linda S Greene, Lolita Buckner Inniss, Bridget J Crawford, Mehrsa Baradaran, Noa Ben-Asher, I Bennett Capers, Osamudia R James, and Keisha Lindsay. Talking About Black Lives Matter and #MeToo. *Wis. JL Gender, & Soc'y*, 34:109, 2019.

[46] Gabriel Grill. Future Protest Made Risky: Examining Social Media Based Civil Unrest Prediction Research and Products. *Computer Supported Cooperative Work (CSCW)*, 30(5-6):811–839, 2021.

[47] Josh Guberman. #ActuallyAutistic Twitter as a Site for Epistemic Resistance. *ACM Transactions on Computer-Human Interaction*, 2022. ISSN 1073-0516. Just Accepted.

[48] Paridhi Gupta, Adrian Petterson, Divyani Motla, and Priyank Chandra. Ladange, Adange, Jeetange: The Farmers' Movement and Its Virtual Spaces. *Proceedings of the ACM on Human-Computer Interaction*, 6(CSCW2):1–34, 2022.

[49] William L Hamilton, Jure Leskovec, and Dan Jurafsky. Diachronic Word Embeddings Reveal Statistical Laws of Semantic Change. *arXiv preprint arXiv:1605.09096*, 2016.

[50] Philip N Howard, Aiden Duffy, Deen Freelon, Muzammil M Hussain, Will Mari, and Marwa Maziad. Opening Closed Regimes: What Was the Role of Social Media During the Arab Spring? *Social Science Research Network*, 2011.

[51] Elke Ichau, Thomas Frissen, and Leen d'Haenens. From #selfie to #edgy. Hashtag networks and images associated with the hashtag #jews on Instagram. *Telematics and Informatics*, 44:101275, 2019.

[52] Sobia Ilyas. The Semiotic Use of Language in the #MeToo Movement: A Multimodal Discourse Analysis. *Hayatian Journal of Linguistics and Literature*, 5(1), 2021.

[53] The Tribune India. Youth trend hashtags in support of farmers. *Tribune*, December 22 2021. URL https://www.tribuneindia.com/news/punjab/youth-trend-hashtags-in-support-of-farmers-180202. Accessed May 9, 2023.

[54] Naveed Iqbal. From tractor to Twitter and Trolley Times: Here's how some Twitter handles took farmers from fields to social media. *Indian Express*, 2021. URL https://indianexpress.com/article/cities/ludhiana/from-tractor-2-

twitter-and-trolley-times-heres-how-some-twitter-handles-took-farmers-from-fields-to-social-media-7630836/. Accessed May 9, 2023.

[55] Md Saiful Islam, Tonmoy Sarkar, Sazzad Hossain Khan, Abu-Hena Mostofa Kamal, SM Murshid Hasan, Alamgir Kabir, Dalia Yeasmin, Mohammad Ariful Islam, Kamal Ibne Amin Chowdhury, Kazi Selim Anwar, et al. COVID-19–Related Infodemic and Its Impact on Public Health: A Global Social Media Analysis. *The American journal of tropical medicine and hygiene*, 103(4):1621, 2020.

[56] Sarah Jaffe. The Collective Power of #MeToo. *Dissent*, 65(2):80–87, 2018.

[57] Shagun Jhaver, Sucheta Ghoshal, Amy Bruckman, and Eric Gilbert. Online Harassment and Content Moderation: The Case of Blocklists. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 25(2):1–33, 2018.

[58] Jeongmin Jo, Sangyeon Kim, and Sangwon Lee. Read It Now, Not Later! Integrating Semantic Recommendation into In-Browser Web Page Scrap Tool. In *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems*, pages 1–9, 2023.

[59] Jeff Johnson, Matthijs Douze, and Hervé Jégou. Billion-Scale Similarity Search with GPUs. *IEEE Transactions on Big Data*, 7(3):535–547, 2019.

[60] Aditya Kamble, Prathamesh Badgujar, Anuj Kadam, Dhruv Shah, and AJ Kadam. Stance Prediction of Tweets on Farmers Protests in India. *International Journal for Research in Applied Science Engineering Technology (IJRASET)*, 10, 2022.

[61] Os Keyes, Chandler May, and Annabelle Carrell. You Keep Using That Word: Ways of Thinking about Gender in Computing Research. *Proceedings of the ACM on human-computer interaction*, 5(CSCW1):1–23, 2021.

[62] Seema Khan, E Combaz, and FE McAslan. Social Exclusion. *Birmingham: GSDRC, University of Birmingham*, 2015.

[63] Iris Kim. 250 million farmers in India are taking a stand against the government as part of the biggest protest in history. Here's why the US and the world should pay attention. *Business Insider*, 2021. URL https://www.businessinsider.com/indian-farmer-strike-largest-protest-history-us-world-pay-attention-2020-12. Accessed March 23, 2022.

[64] Myeong Gyu Kim, Minjung Kim, Jae Hyun Kim, and Kyungim Kim. Fine-Tuning BERT Models to Classify Misinformation on Garlic and COVID-19 on Twitter. *International Journal of Environmental Research and Public Health*, 19(9):5126, 2022.

[65] Colin Klein, Ritsaart Reimann, Ignacio Ojea Quintana, Marc Cheong, Marinus Ferreira, and Mark Alfano. Attention and counter-framing in the Black Lives Matter movement on Twitter. *Humanities and Social Sciences Communications*, 9(1):1–12, 2022.

[66] Rachel Kuo. Racial justice activist hashtags: Counterpublics and discourse circulation. *New Media & Society*, 20(2):495–514, 2018.

[67] Andrey Kutuzov, Lilja Øvrelid, Terrence Szymanski, and Erik Velldal. Diachronic word embeddings and semantic shifts: a survey. *arXiv preprint arXiv:1806.03537*, 2018.

[68] Mirko Lai, Viviana Patti, Giancarlo Ruffo, and Paolo Rosso. Stance Evolution and Twitter Interactions in an Italian Political Debate. In *Natural Language Processing and Information Systems: 23rd International Conference on Applications of Natural Language to Information Systems, NLDB 2018, Paris, France, June 13-15, 2018, Proceedings 23*, pages 15–27. Springer, 2018.

[69] Mirko Lai, Marcella Tambuscio, Viviana Patti, Giancarlo Ruffo, and Paolo Rosso. Stance polarity in political debates: A diachronic perspective of network homophily and conversations on Twitter. *Data & Knowledge Engineering*, 124:101738, 2019.

[70] Mark Landler. Trump, Saying 'Mere Allegation' Ruins Lives, Appears to Doubt #MeToo Movement. *The New York Times*, 10, 2018. URL https://www.nytimes.com/2018/02/10/us/politics/trump-porter-me-too-movement.html. Accessed April 18, 2023.

[71] Jun Lang, Wesley W Erickson, and Zhuo Jing-Schmidt. #MaskOn! #MaskOff! Digital polarization of mask-wearing in the United States during COVID-19. *PloS one*, 16(4): e0250817, 2021.

[72] Calvin A Liang, Sean A Munson, and Julie A Kientz. Embracing Four Tensions in Human-Computer Interaction Research with Marginalized People. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 28(2):1–47, 2021.

[73] Sheng-Chieh Lin and Jimmy Lin. A Dense Representation Framework for Lexical and Semantic Matching. *ACM Transactions on Information Systems*, 41(4):1–29, 2023.

[74] Lisa Lindqvist and Simon Lindgren. Mapping an Emerging Hashtag Ecosystem: Connective Action and Interpretive Frames in the Swedish #MeToo Movement. *Feminist Media Studies*, pages 1–18, 2022.

[75] Suman Kalyan Maity, Ritvik Saraf, and Animesh Mukherjee. #Bieber+ #Blast= #Bieberblast: Early Prediction of Popular Hashtag Compounds. In *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing*, pages 50–63, 2016.

[76] Matej Martinc, Petra Kralj Novak, and Senja Pollak. Leveraging Contextual Embeddings for Detecting Diachronic Semantic Shift. *arXiv preprint arXiv:1912.01072*, 2019.

[77] Mary L McHugh. Interrater reliability: the kappa statistic. *Biochemia medica*, 22(3): 276–282, 2012.

[78] Cristina Menghini, Justin Uhr, Shahrzad Haddadan, Ashley Champagne, Bjorn Sandstede, and Sohini Ramachandran. The Drift of #MyBodyMyChoice Discourse on Twitter. In *14th ACM Web Science Conference 2022*, pages 110–117, 2022.

[79] Dibyendu Mishra, Syeda Zainab Akbar, Arshia Arya, Saloni Dash, Rynaa Grover, and Joyojeet Pal. Rihanna versus Bollywood: Twitter Influencers and the Indian Farmers' Protest. *arXiv preprint arXiv:2102.04031*, 2021.

[80] Stein Monteiro. Farmer protests in India and the mobilization of the online diaspora on Twitter. *Social Science Research Network*, 2021.

[81] Rachel Elizabeth Moran, Stephen Prochaska, Isabelle Schlegel, Emelia May Hughes, and Owen Prout. MISINFORMATION OR ACTIVISM: MAPPING NETWORKED MORAL PANIC THROUGH AN ANALYSIS OF #SAVETHECHILDREN. *AoIR Selected Papers of Internet Research*, 2021.

[82] Aaron Mueller, Zach Wood-Doughty, Silvio Amir, Mark Dredze, and Alicia Lynn Nobles. Demographic Representation and Collective Storytelling in the Me Too Twitter Hashtag Activism Movement. *Proceedings of the ACM on human-computer interaction*, 5(CSCW1):1–28, 2021.

[83] Diana C Mutz. *Hearing the Other Side: Deliberative Versus Participatory Democracy.* Cambridge University Press, 2006.

[84] Behzad Naderalvojoud, Ahmet Selman Bozkir, and Ebru Akcapinar Sezer. INVESTIGATION OF TERM WEIGHTING SCHEMES IN CLASSIFICATION OF IMBALANCED TEXTS. In *Lisbon: Proceedings of European Conference on Data Mining (ECDM)*, pages 39–46, 2014.

[85] Kumari Neha, Vibhu Agrawal, Vishwesh Kumar, Tushar Mohan, Abhishek Chopra, Arun Balaji Buduru, Rajesh Sharma, and Ponnurangam Kumaraguru. A Tale of Two Sides: Study of Protesters and Counter-protesters on #CitizenshipAmendmentAct Campaign on Twitter. In *14th ACM Web Science Conference 2022*, pages 279–289, 2022.

[86] Ashwin Sanjay Neogi, Kirti Anilkumar Garg, Ram Krishn Mishra, and Yogesh K Dwivedi. Sentiment analysis and classification of Indian farmers' protest using Twitter data. *International Journal of Information Management Data Insights*, 1(2):100019, 2021.

[87] Marco Polignano, Valerio Basile, Pierpaolo Basile, Giuliano Gabrieli, Marco Vassallo, and Cristina Bosco. A hybrid lexicon-based and neural approach for explainable polarity detection. *Information Processing & Management*, 59(5):103058, 2022.

[88] Anna Priyadarshini. What happened to the woman allegedly gangraped at Tikri?, 2021. URL https://www.newslaundry.com/2021/05/15/what-happened-to-the-woman-allegedly-gangraped-at-tikri. Accessed April 5, 2023.

[89] Yashwant Raj. US Congress report warns of repercussions over farmer protests in India. *Hindustan Times*, 2021. URL https://www.hindustantimes.com/world-news/us-congress-report-warns-of-repercussions-over-farmer-protests-in-india-101615869352898.html. Accessed December 7, 2021.

[90] Kunal Rao and Gaurav Dixit. Genuineness of Social media Support: Role of languages and influencers in the Case of Farmers' Protest in India. In *PACIS 2022 Proceedings*, page 248, 2022.

[91] Rashawn Ray, Melissa Brown, Neil Fraistat, and Edward Summers. Ferguson and the death of Michael Brown on Twitter: #BlackLivesMatter, #TCOT, and the evolution of collective identities. *Ethnic and racial studies*, 40(11):1797–1813, 2017.

[92] Ricardo Rei, Nuno Miguel Guerreiro, and Fernando Batista. Automatic Truecasing of Video Subtitles Using BERT: A Multilingual Adaptable Approach. In *Information Processing and Management of Uncertainty in Knowledge-Based Systems: 18th International Conference, IPMU 2020, Lisbon, Portugal, June 15–19, 2020, Proceedings, Part I 18*, pages 708–721. Springer, 2020.

[93] Nils Reimers and Iryna Gurevych. Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3982–3992, 2019.

[94] Julio CS Reis, Philipe Melo, Kiran Garimella, Jussara M Almeida, Dean Eckles, and Fabrício Benevenuto. A Dataset of Fact-Checked Images Shared on WhatsApp During the Brazilian and Indian Elections. In *Proceedings of the international AAAI conference on web and social media*, volume 14, pages 903–908, 2020.

[95] Reuters. Timeline: Indian farmers' protests against agricultural laws. *Reuters*, 2021. URL https://www.reuters.com/markets/commodities/indian-farmers-protests-against-agricultural-laws-2021-11-19/. Accessed April 9, 2023.

[96] Eugenia Ha Rim Rho. *The Design of Online Environments (Political Hashtags) and the Quality of Democratic Discourse At-Scale*. University of California, Irvine, 2020.

[97] Eugenia Ha Rim Rho and Melissa Mazmanian. Hashtag Burnout? A Control Experiment Investigating How Political Hashtags Shape Reactions to News Content. *Proceedings of the ACM on human-computer interaction*, 3(CSCW):1–25, 2019.

[98] Eugenia Ha Rim Rho, Gloria Mark, and Melissa Mazmanian. Fostering Civil Discourse Online: Linguistic Behavior in Comments of #MeToo Articles across Political Perspectives. *Proceedings of the ACM on human-computer interaction*, 2(CSCW): 1–28, 2018.

[99] Andrew S Ross. Discursive delegitimisation in metaphorical #secondcivilwarletters: an analysis of a collective Twitter hashtag response. *Critical Discourse Studies*, 17(5): 510–526, 2020.

[100] Koustuv Saha, Sang Chan Kim, Manikanta D Reddy, Albert J Carter, Eva Sharma, Oliver L Haimson, and Munmun De Choudhury. The Language of LGBTQ+ Minority Stress Experiences on Social Media. *Proceedings of the ACM on human-computer interaction*, 3(CSCW):1–22, 2019.

[101] C Sahithi, Y Sreeja, S Akhil, K Taruni, and CC Sobin. Twitter Sentiment Analysis Using Machine Learning Techniques: A Case Study Based on Farmers Protest. In *Machine Learning in Information and Communication Technology: Proceedings of ICICT 2021, SMIT*, pages 187–196. Springer, 2022.

[102] Mila Samdub. Governance Interface: The Index in Modi's India. *Public Culture*, 34 (3):473–493, 2022.

[103] Abraham C Sanders, Rachael C White, Lauren S Severson, Rufeng Ma, Richard McQueen, Haniel C Alcântara Paulo, Yucheng Zhang, John S Erickson, and Kristin P Bennett. Unmasking the conversation on masks: Natural language processing for topical sentiment analysis of COVID-19 Twitter discourse. *AMIA Summits on Translational Science Proceedings*, 2021:555, 2021.

[104] Maarten Sap, Dallas Card, Saadia Gabriel, Yejin Choi, and Noah A Smith. The Risk of Racial Bias in Hate Speech Detection. In *Proceedings of the 57th annual meeting of the association for computational linguistics*, pages 1668–1678, 2019.

[105] Dietram A Scheufele and David Tewksbury. Framing, Agenda Setting, and Priming: The Evolution of Three Media Effects Models. *Journal of communication*, 57(1):9–20, 2007.

[106] Vipul Shah, Amey Bhole, Sandeep S Udmale, and Vijay Sambhe. A Deep Multi-kernel Uniform Capsule Approach for Hate Speech Detection. In *Distributed Computing and Intelligent Technology: 18th International Conference, ICDCIT 2022, Bhubaneswar, India, January 19–23, 2022, Proceedings*, pages 265–271. Springer, 2022.

[107] Divya Siddarth and Joyojeet Pal. Engaging the Crowd: Social Movement Building via Online Bystander Mobilization. In *Proceedings of the 2020 International Conference on Information and Communication Technologies and Development*, pages 1–13, 2020.

[108] Karin Skill, Sergio Passero, and Marie Francisco. Assembling Amazon Fires through English Hashtags. Materializing Environmental Activism within Twitter Networks. *Computer Supported Cooperative Work (CSCW)*, 30(5-6):715–732, 2021.

[109] Kate Starbird, Ahmer Arif, and Tom Wilson. Disinformation as Collaborative Work: Surfacing the Participatory Nature of Strategic Information Operations. *Proceedings of the ACM on Human-Computer Interaction*, 3(CSCW):1–26, 2019.

[110] Leo Graiden Stewart, Ahmer Arif, A Conrad Nied, Emma S Spiro, and Kate Starbird. Drawing the Lines of Contention: Networked Frame Contests Within #BlackLivesMatter Discourse. *Proceedings of the ACM on Human-Computer Interaction*, 1(CSCW): 1–23, 2017.

[111] Mirac Suzgun, Stuart M Shieber, and Dan Jurafsky. string2string: A Modern Python Library for String-to-String Algorithms. *arXiv preprint arXiv:2304.14395*, 2023.

[112] Shiliang Tang, Qingyun Liu, Megan McQueen, Scott Counts, Apurv Jain, Haitao Zheng, and Ben Y Zhao. Echo Chambers in Investment Discussion Boards. In *Eleventh International AAAI Conference on Web and Social Media*, 2017.

[113] Nandan Thakur, Nils Reimers, Johannes Daxenberger, and Iryna Gurevych. Augmented SBERT: Data Augmentation Method for Improving Bi-Encoders for Pairwise Sentence Scoring Tasks. *arXiv preprint arXiv:2010.08240*, 2020.

[114] Spencer Torene, Andrew Follmann, Thomas Teague, Peter Chang, and Blake Howald. Automated Hashtag Hierarchy Generation Using Community Detection and the Shannon Diversity Index, with Applications to Twitter and Parler. *International Journal of Semantic Computing*, 16(04):473–496, 2022.

[115] Ralph H Turner. The Public Perception of Protest. *American sociological review*, pages 815–831, 1969.

[116] Kobie Van Krieken and José Sanders. Diachronic changes in forms and functions of reported discourse in news narratives. *Journal of pragmatics*, 91:45–59, 2016.

[117] Courtland VanDam and Pang-Ning Tan. Detecting Hashtag Hijacking from Twitter. In *Proceedings of the 8th ACM Conference on Web Science*, pages 370–371, 2016.

[118] Ashwani Varshney, Yatin Kapoor, Anjali Thukral, Richa Sharma, and Punam Bedi. Performing Sentiment Analysis on Twitter Data Using Deep Learning Models: A Comparative Study. In *Advances in Data and Information Sciences: Proceedings of ICDIS 2021*, pages 371–381. Springer, 2022.

[119] James P Walsh. Social media and moral panics: Assessing the effects of technological change on societal reaction. *International Journal of Cultural Studies*, 23(6):840–859, 2020.

[120] Wenhui Wang, Furu Wei, Li Dong, Hangbo Bao, Nan Yang, and Ming Zhou. MiniLM: Deep Self-Attention Distillation for Task-Agnostic Compression of Pre-Trained Transformers. *Advances in Neural Information Processing Systems*, 33:5776–5788, 2020.

[121] Guobin Yang. Narrative Agency in Hashtag Activism: The Case of #BlackLivesMatter. *Media and communication*, 4(4):13, 2016.

[122] Michele Zappavigna. *Searchable Talk: Hashtags and Social Media Metadiscourse*. Bloomsbury Publishing, 2018.

Appendices

Appendix A

Tables to Assess for Causality

Time Series	ADF	Threshold (1%)	Threshold (5%)	Conclusion
Semantic polarization	-9.834	-3.497	-2.89	stationary
Daily count of offspring hashtags	-7.434	-3.497	-2.89	stationary
Daily count of newly emerged offspring hashtags	-7.030	-3.49	-2.86	stationary

Table A.1: ADF Test to Check Stationarity for H1 and H2.

Time Series	ADF	Threshold (1%)	Threshold (5%)	Conclusion
Semantic polarization	-9.5	-3.58	-2.89	stationary
Volume of tweets that exclusively use parent hashtag	-8.3	-3.58	-2.89	stationary

Table A.2: ADF Test to Check Stationarity for H3.

Appendix B

Data Collection and Classification Using the Parent Hashtag

We initially collected tweets from September 9, 2020, to November 11 November 2021, using the parent hashtag "#FarmersProtest" as our primary filter. We then refined our dataset by including only tweets that contained the token "farmer" without any additional offspring hashtags. To ensure the consistency and reliability of our analysis, we considered only tweets in English originating from the Indian subcontinent. We followed similar pre-processing steps as mentioned in Chapter 3.

Our classification process began with manually labeling 100 samples of tweets each from protest and counter-protest groups based on our contextual knowledge and understanding of the data. Using these manually labeled samples, we fine-tuned a BERT-based classification model to classify the remaining unlabeled tweets into two sub-groups: protest (pro-farmer) and counter-protest (anti-farmer).

Our final dataset consists of 7,165 tweets, of which 6,566 were pro-farmer (protest), and 609 were anti-farmer (counter-protest). The results from the classification model and the overall split of data are depicted in Table B.1 and Table B.2, respectively.

Model	Accuracy	Precision	Recall	F1 Score
BERT (fine-tuned)	0.973	0.916	0.77	0.84

Table B.1: Classification Metrics.

Descriptive Statistics	Total	Protest	Counter-Protest
Mean length of tweets	26	25	32
Median length of tweets	25	23	42
25th percentile of tweet length	14	14	24
75th percentile of tweet length	38	37	42

Table B.2: Descriptive Statistics of Tweets with Parent Hashtags Identified as Protest and Counter-Protest.

B.1 Measuring Semantic Polarization Shift in Protests and Counter-Protests which Use Parent Hashtags

We calculate the semantic polarization around the token "farmer" between protests and counter-protests that exclusively use the parent hashtag. The technique we use to determine semantic polarization is identical to the one detailed in [31]. The daily average of semantic polarity between protests and counter-protests that exclusively use the parent hashtag is observed to be 0.54, 37% less than the semantic polarity observed between protests and counter-protests that use offspring hashtags (see Table B.3). Figure B.1 showcases the shift in semantic polarization between protests and counter-protests that exclusively use the parent hashtag throughout the protest.

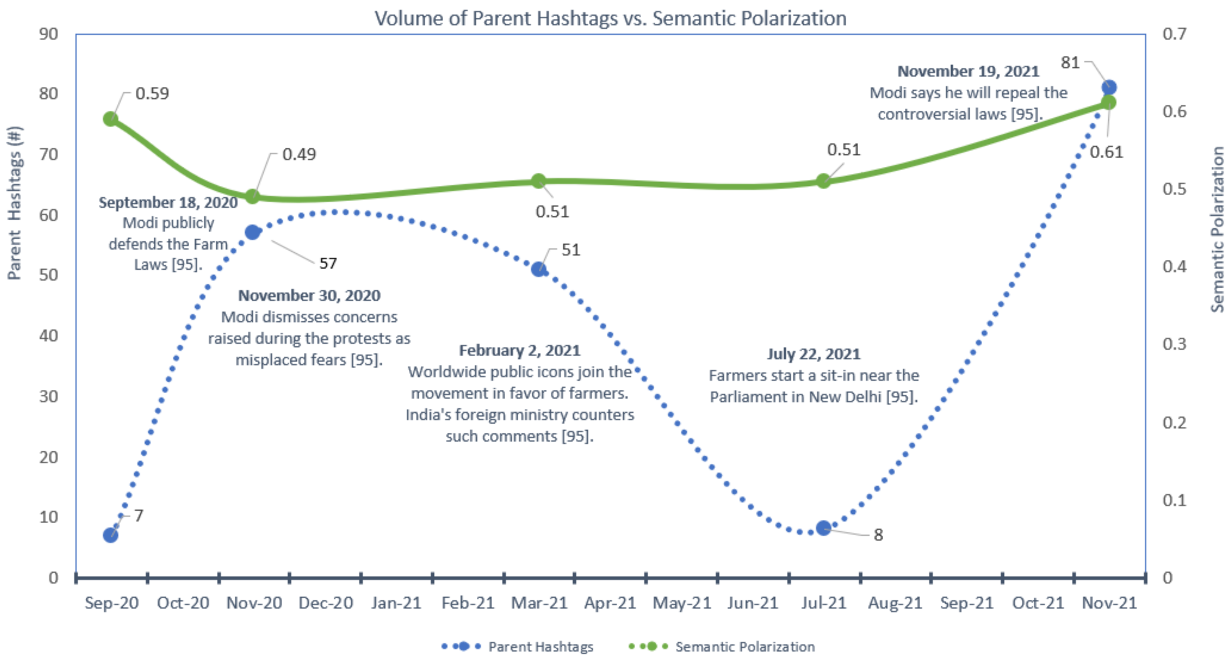

Figure B.1: Evolution of Semantic Polarity Around the Word "farmer" Between Protests and Counter-Protests That Exclusively Use the Parent Hashtag.

Daily Average (Only Parent Hashtag)	Daily Average (Offspring Hashtags)
0.54	0.74

Table B.3: Daily Average in Semantic Polarization Around the Token "farmer".