#BlackLivesMatter: A Network Analysis

Ray Crist
Honors Thesis

Professor Kimberly Rogers, Advisor
Department of Sociology
Dartmouth College
June 6, 2022

ACKNOWLEDGMENTS

Many people contributed in ways large and small to the completion of this thesis.

First, I’d like to express my deepest gratitude to Professor Kimberly Rogers for her immense support as my thesis advisor. Ever since I took her Introductory Sociology course two years ago, Professor Rogers has been a mentor to me, and over the last year, Professor Rogers has been a vital resource. Through frequent meetings and constant encouragement, Professor Rogers gave me the resources to complete this thesis.

I’d additionally like to express special thanks to Professors Shaonta' Allen, Peter Mucha, Sunmin Kim, and Antonio Sirianni for their invaluable advice and feedback, without which my finished product would have been lacking.

Finally, I’d like to thank my family and friends, many of whom have been invested in this project since the beginning. I am extremely grateful for their kind words and support throughout this project.

ABSTRACT

In the aftermath of George Floyd's murder, Twitter experienced one of its largest surges in activity ever, as users expressed their outrage and signaled their support for Black Lives Matter. Given the growing importance of online activism, understanding the diffusion of #BlackLivesMatter on Twitter during this period is critical for understanding modern social movements. This thesis examines hashtag diffusion on Twitter by undertaking a network analysis. First, I broadly characterize adopter behavior during the week after Floyd's death, finding that the earliest adopters tended to be well-connected on Twitter and that many users outside the United States expressed their support for the movement. Using geotagged tweets to match users to county subdivisions, I then examine how hashtag adoption behavior varies based on county subdivision characteristics including population density, income, racial demographics, and rates of police violence. I find that the diffusion of #BlackLivesMatter on Twitter was a predominantly urban phenomena, driven by adopters from large cities and areas with large Black populations, as well as by adopters from Minneapolis. Finally, I use probabilistic models of social contagion to test whether #BlackLivesMatter represents a social contagion and find that homophily and factors external to Twitter best explain hashtag adoption over the study period, but that models of contagion best explain the first forty-eight hours of adoption, with users from more urban and diverse environments requiring less social reinforcement to adopt the hashtag.

TABLE OF CONTENTS

CHAPTER 1: INTRODUCTION

Social media sites have often been described as transformational for the modern social movement (Borge-Holthoefer 2014, Tufekci 2014). In the United States, The Black Lives Matter movement has been intimately associated with social media, ever since the phrase “Black Lives Matter” was invented by Alicia Garza on Facebook in the aftermath of the George Zimmerman verdict and subsequently turned into a hashtag by Patrisse Cullors (Brown 2015).

Evidence has accumulated about how Black Lives Matter activists have used social media to counteract negative media messages about unrest (Welles and Jackson 2019), build community among movement organizers (Edrington and Lee 2018), and build organizational capacity (Mundt, Ross, and Burnett 2018). Furthermore, #BlackLivesMatter activity on Twitter has been linked to increased next-day protest activity (De Choudhury et al. 2016) and responses to the movement from political elites (Freelon, McIlwain, and Clark 2018).

Black Lives Matter has steadily grown over the last several years, but the death of George Floyd in May 2020 marked a turning point for the movement. The Twitter platform witnessed one of its busiest periods ever as users reacted to Floyd’s death and expressed support for the movement (Wu et al. 2021), and the succeeding protests have been described as potentially the largest social movement in American history (Buchanan, Bui and Patel 2020). Understanding how the George Floyd moment spread online has major sociological implications, and I undertake a networks-based methodology towards this end.

Social network analysis has become a critical feature of the social sciences ever since Granovetter's (1973) pioneering work on the strength of weak ties demonstrated the role of networks in transmitting novel information. Later scholars have carried on Granovetter's work via the basic model of contagion, in which an entity (such as a behavior, belief, or use of innovation) spreads through contact with "infected" individuals (Bakshy et al. 2012). With the rise of social media making huge troves of social network data available, and the rise of computing power facilitating complex forms of modelling, this has become a popular field of study (Hanappi 2017, Himelboim 2017, Guilbeault, Becker, and Centola 2018).

In the following thesis, I examine the characteristics of the spread of #BlackLivesMatter on Twitter in the week after George Floyd's murder from both a geospatial and a networks perspective. I find that hashtag adoption varied based upon region, with users from more urban and diverse county subdivisions more likely to adopt #BlackLivesMatter. I additionally find that adopters from Minnesota and from predominantly Black areas were important early adopters of the hashtag. Finally, I find that online social contagion is critical for the initial diffusion of #BlackLivesMatter, while later adoption can be explained by factors external to Twitter.

My thesis is laid out in the following way: In Chapter 2, I examine all relevant literature, first examining the relevant social movement literature in order to understand why hashtag adoption matters, then examining the networks and contagion literature in order to understand how hashtag adoption occurs. In Chapter 3, I outline my three research questions. Each of the following Chapters 4 through 6 are dedicated towards answering one of these questions. In Chapter 4, I examine the overall characteristics of hashtag adoption. In Chapter 5, I examine how adoption behavior varied based upon geographic characteristics. Finally, in Chapter 6, I test whether hashtag adoption aligns with models of social contagion. Given that my methodology for each research question builds upon my findings for the prior research question, I begin each Chapter 4 through 6 by defining my methods for the relevant research question, then proceed by describing my results. I conclude with Chapter 7, a brief summary of my findings and how they are significant.

CHAPTER 2: LITERATURE REVIEW

In the following chapter, I analyze literature relevant to the diffusion of #BlackLivesMatter on Twitter in two parts:

First, I examine literature specific to social movements and Black Lives Matter in order to clarify how online activism functions to further movement goals. I begin with a brief history of Black Lives Matter and its evolving online presence. Then I identify the challenges of contextualizing online activism within traditional social movement paradigms, and identify the framework of connective action, which highlights the role of the network in movement growth, as a useful lens for interpreting online activism. Actor-network theory illuminates the role of digital technologies in shaping the capacities of online activism to affect change; thus, I analyze how the specific characteristics of social media platforms and hashtags shape the outcomes of online activism. I conclude the first section of my literature review with an examination of theories of the public sphere, finding that online activism is a powerful generator of counterpublics and a tool for shaping discourse.

Next, I examine literature on social networks, first broadly and then more specifically in regards to theories of social contagion (social phenomena which transmit across networks). While not specific to social movements, the social networks literature provides us with the tools to understand how connective action spreads. I identify threshold models as a valuable tool for understanding how individuals decide to engage in socially-risky behavior, then expand on this notion with the network-specific models of simple and complex contagion. After examining empirical studies of contagion and their limitations, I conclude this chapter by articulating how my research expands upon previous work.

#BlackLivesMatter, Movements, and the Public Sphere

The Digital Roots of Black Lives Matter

The Black Lives Matter movement has been intimately associated with social media since its inception. On the day of George Zimmerman's acquittal for the killing of Trayvon Martin, July 13th 2013, activist Alicia Garza coined the eponymous slogan in a series of Facebook posts jointly titled "A Love Letter to Black People". Patricia Cullors, Garza's friend and fellow activist, converted the phrase into a hashtag, and soon the duo, along with fellow activist Ayo (Opal) Tometi, began attempting to organize around #BlackLivesMatter (Cobb 2016), starting by creating a social media group in which Black activists could "tell [their] stories, share grief, share rage, [and] collaborate together" (Garza and Kauffman 2015).

It would not be until a year later, when Michael Brown was shot by a white police officer in Ferguson, Missouri that #BlackLivesMatter would reach a large audience. The Ferguson protests of August 2014 have been described as the birth of the new racial equality movement (Chase 2017), and social media played a critical role in elevating the protests in Ferguson to a national stage. Tufecki and Wilson (2012) articulated the role of "citizen journalists" in social movements -- individuals who document protest events on social media and thus transmit information to a wide audience, evading censors. During the Ferguson protests, Twitter users like Johnetta Elias fulfilled this role, amassing large followings by sharing on-the-ground details, exposing the wrongdoings of police officers and contradicting mainstream media narratives, which failed to recognize the problem of police violence (Jackson and Welles 2016). Many of the users who rose to prominence in the aftermath of Ferguson were not elites, nor were many affiliated with Garza, Cullors, and Tometi; instead, they represented a loose collection of Ferguson citizens, Black activists, and non-Black allies who rallied around a shared set of hashtags (Freelon et al. 2016).

In the aftermath of Ferguson, the movement's messaging had not yet crystallized, evidenced by the fact that #BlackLivesMatter did not breach the top ten hashtags used in reference to the protests. Instead, hashtags like #Ferguson and #MikeBrown dominated Twitter (Freelon et al. 2016), which might have suggested that the Ferguson protests would be an isolated incident. This would not be the case. In early September, Cullors and fellow activist Darnell Moore organized Black Lives Matter Freedom Rides to Ferguson (Solomon 2014; Moore and Cullors 2014). Garza has noted that the Freedom Rides were a crucial turning point in the movement's history, as they concluded with a group agreement for future coordinated actions (Garza and Kauffman 2015). Thus, Ferguson was instrumental in fostering the coalition of formal organizations and activists that would come to lead the Black Lives Matter movement. This coalition would be crucial for organizing protest events and raising awareness to allow #BlackLivesMatter to reach a massive audience.

The winter of 2014 witnessed the completion of #BlackLivesMatter's evolution from a little-known phrase to the symbol of the new movement. During the fall, the hashtag had slowly gained momentum among activist circles. When a grand jury chose not to indict Derek Wilson for the killing of Michael Brown on November 24th, the hashtag reached critical mass. Tweets including #BlackLivesMatter rose from just over two thousand the previous day to over one-hundred thousand (Freelon et al. 2016). Unlike previous bursts of attention, the most retweeted users on November 24th were celebrities, writers, and ordinary Twitter users. Even users which generally did not engage in political issues, such as humor accounts, chimed in (Freelon et al. 2016). From 2015 on, #BlackLivesMatter was regularly included in hundreds of thousands of tweets during spikes of hashtag usage, often following police killings and major court developments. Black users were often the users most engaged with #BlackLivesMatter, but during spikes, a more diverse group of users engaged (Olteanu, Weber, and Gatica-Perez 2015).

Even in the earliest stages of the movement, social media was a powerful force for strengthening the movement. During the Ferguson protests, social media was highly effective at nationalizing a local issue and elevating marginalized narratives to a wider audience. Furthermore, the use of social media facilitated the ability of hundreds of thousands of people to signal their opinion regarding incidents of police brutality without a major organizational apparatus in-place. The usage of hashtags was often invaluable in consolidating conversations across disparate groups.

At the same time, the early experience of Black Lives Matter points to the challenge researchers face when studying social media in social movements -- social media users are often only weakly linked to others in the movements they support. During the Ferguson protests, many of the most prominent voices online, such as Johnetta Elias, had no connection to the coalition being formed by Garza, Tometi, and Cullors, and have been critical of those who attribute the birth of Black Lives Matter to them (Cobb 2016). It feels obvious that users such as Elias who provided real-time documentation of Ferguson protests would be considered a part of Black Lives Matter, but it is less clear that the pop-culture figures and entertainers who tweeted #BlackLivesMatter in the aftermath of Derek Wilson’s non-indictment should be considered movement members, or if they would even identify as such. As the movement grew, such questions about membership would become increasingly common, as more users expressed support online. Since my research is specifically examining the spread of #BlackLivesMatter, it is important to tackle the question of what it means to engage with a social movement online, if my findings are to have any significance.

Defining Online Social Movements

Scholars have struggled to articulate how the diversity of online social movement activity fits into paradigms of social movement scholarship. Diani (1992:1), for example, defines a social movement as "consisting in networks of informal interaction between a plurality of individuals, groups, and/or organizations, engaged in a political and/or cultural conflict, on the basis of a shared collective identity". While Black Lives Matter may be readily described as a network of interactions, it is less obvious what, if any, collective identity unites movement participants. Does a Twitter user who tweets using #BlackLivesMatter, for example, consider themselves as a Black Lives Matter activist? Some authors have handled this question by arguing that signaling support for a movement online is not sufficient to be considered a movement participant. Sending a tweet does not convey the same level of commitment that standing in the streets under tear gas does. This is the central criticism implied by describing hashtag activism as *slacktivism* -- existing across weak-tie networks, it lacks the requirement of commitment that other forms of activism do (Gladwell 2010). A large body of social movement research has highlighted the importance of high participant commitment and/or a cohesive sense of identity in a movement's capacity to affect change (McAdam 1986, Olson 1965, Melucci 1995); without these properties, it can be hard to decipher the value of hashtag activism. Some studies have solved this problem by explicitly limiting their focus to online engagement which motivates individuals to attend protests (Jost et al. 2018, Tufekci and Wilson 2012, Greijdanus et al. 2020). Yet hashtag activism is increasingly becoming a major feature of modern social movements; overlooking such behavior risks overlooking what this behavior means. In order to understand such behavior, and how it affects the overall Black Lives Matter movement, we need a new conceptual framework.

Logics of Connective and Collective Action

Hashtag activism can be understood as a form of connective action, characterized by low-commitment actions that emphasize the development of an individual's identity rather than the adoption of a collective identity. Movements can be characterized by the strategies they use to solve the central dilemma of organizing: how to get individuals to engage, given that they will benefit from the products of a successful movement regardless of participation. Twentieth-century social movements typically incentivized participation through what Olson (1965) defined as the logic of collective action: the development of a strong collective identity and formal organizations to manage participation. By contrast, Bennett and Sederberg (2012:748) observe that many modern movements follow an emergent "logic of connective action", in which users participate through acts of personal expression across digitally mediated, "large-scale, fluid social networks". Rather than generating a strong sense of collective identity, connective action emphasizes the development of individual identity through the use of personal action frames, in which the user adapts movement messaging to relate to their own experience. The use of personal action frames provides opportunity for recognition and validation by one's peers, thus motivating individuals to propagate messages through the network.

The network replaces the formal organization as the central structure in connective action. Bennet and Sederberg's (2012) definition of connective action relies on actor-network theory to explain how digital technologies make this possible. Actor-network theory, developed by Latour (2005), Law (1992), and others, states that social systems can be best understood as networks of diverse actors, human and non-human, which collude to produce effects which can be attributed to the overall system. Within networks of connective action, non-human actors, such as mobile devices and the features of social networking sites, facilitate connective action by reducing the communicative costs of participation (Lupia and Sin 2003) and providing organizational capacity. For example, in previous social movements, an individual may be tasked with keeping track of members and events. In modern movements, group affiliation may be logged automatically based on an individual's affiliation with a Facebook group, and events can be easily shared with large groups of people. When digital technologies replace organizational demands that would otherwise require high commitment from individuals, connective action becomes possible.

Black Lives Matter can be analyzed as a hybrid movement, engaging both the logics of connective and collective action. Organizations such as the Black Lives Matter Network and Movement for Black Lives engage collective action methods to further movement objectives, such as by occupying public spaces, hosting "die-ins", and providing demands to political elites (Rickford 2016). On the other hand, social media users engage connective action methods by sharing movement-related messaging in dialogue with their peers. Alfonzo (2021) examined a sample of tweets from highly-active Black Lives Matter users over the last ten years and found that 52% of users shared a large amount of personal action frames in their tweets, while 92% had shared at least one, where personal action frames can be distinguished by their frequent use of the "I" pronoun and their individualized expressions of emotion, judgment, and solidarity. One complication of Bennett and Sederberg's theory of connective action is that #BlackLivesMatter discourse, despite its decentralized nature, still reflects the development of a collective identity.

Melucci (1995:44) described collective identity as the “process of ‘constructing’ an action system”, where activists agree upon a shared set of norms and behaviors. The invocation of racial identity in #BlackLivesMatter discourse, in addition to facilitating users’ personal action frames, is used to prevent domination of the conversation from non-Black Twitter users, articulate Black users’ moral authority on the issue, and debate acceptable actions (Wilkins, Livingstone, and Levine 2019). Yet affiliating with a collective identity is not a prerequisite for engaging with #BlackLivesMatter, and the widespread use of personal action frames means that each individual is free to choose how they relate to the movement.

Characterizing the usage of #BlackLivesMatter as a form of connective action is analytically useful for several reasons. First, we overcome the slacktivism debate by recognizing that online activism is not a “lesser” form of higher-risk activism. Instead, online activism represents a unique form of activism with its own set of motivating forces and logic. Furthermore, examining online activism through the lens of connective action suggests that in order to understand the spread of #BlackLivesMatter, we must turn our attention to the mechanisms of message propagation which Bennett and Sederberg highlight as central to connective action. But first: if online activism is not reducible to slacktivism, then we still need to specify how online activism contributes to movement goals.

To finally address this point, we will leverage Bennett and Sederberg’s identification of digital media components as critical actors within connective action. Digital technologies shape the capacity of connective action networks to affect change, much like how formal organizations shape the capacity of collective action. By examining the unique properties of the platform (Twitter) and the hashtag (#BlackLivesMatter) in fostering different forms of movement activity, we will clarify how the decentralized propagation of movement messaging online impacts the Black Lives Matter movement.

The Platform: Twitter as a Public Sphere

The social network platform Twitter constitutes part of the "public sphere" in which users can contest social values and the meaning of major events. Twitter has been the epicenter of online Black Lives Matter Activity and is especially relevant when examining a movement for racial justice. The relationship between Black expression and Twitter precedes #BlackLivesMatter; "Black Twitter" stands out as a cultural phenomenon produced by a combination of technological and racial processes, including favorable algorithms, a network structure which facilitates information flow, white fascination, and high early adoption rates of Twitter by Black Americans (Sharma 2013; Brock 2012). More importantly, Twitter's public-facing nature makes it an ideal location for social movements to spread. Twitter's features, including hashtags, retweets, and mentions, make it challenging if not impossible for users to contain their audience (Marwick and boyd 2010). Twitter's branding emphasizes its public nature in contrast to more private, friend-oriented platforms like Facebook; Twitter's About page prominently announces its intention to "serve the public conversation" (Twitter, Inc. n.d.). Bruns and Moe (2014) outline the layers of communication on Twitter which address audiences of distinct sizes and find that hashtag-based exchanges represent the widest level of communication. The use of topical hashtags, they argue, is akin to making "a speech at a public gathering" (Bruns and Moe 2014).

To understand Twitter's role in broader society, it is valuable to consider Habermas' theory of the public sphere. First outlined in The Structural Transformation of the Public Sphere (1989), Habermas' public sphere is a space distinct from the economic sphere and the state, structured by mass media institutions, in which “mediated political communication” is “carried on by an elite” (Habermas 2006:416). Modern scholars have modified Habermas’ theory of the public sphere to fit the contemporary landscape; for example, Bruns and Highfield (2016) argue that while the modern fragmented media landscape has rendered a unified public sphere impossible, the public sphere is now composed of a series of “public sphericules” which contain subsets of the discourse. Furthermore, they argue, Twitter is one such sphericule.

While some scholars break up the public sphere according to technological and platform barriers, a separate set of scholars has examined the public sphere as splintering along the lines of ideologies. Fraser (1990) argues that the notion of a unified public sphere is not only unrealistic, but that a unified sphere cannot exist in practice in a stratified society, because the ability to set the discourse becomes yet another method of domination and distinction. Instead, Fraser argues that the existence of counterpublics -- publics composed of marginalized groups which contest hegemonic discourse -- suggests an alternative public sphere, which she theorizes as a structured space in which negotiation and ideological combat between publics can occur. The almost instantaneous rise of #AllLivesMatter in response to #BlackLivesMatter’s first major viral event in December 2014 is an ideal example of publics engaging in ideological conflict on Twitter; during this episode, the hashtags #BlackLivesMatter and #AllLivesMatter became “reified signs” which emblemized the debate over society’s valuation of “Black male lives in relation to ‘all’ … lives” (Carney 2016:194). The involvement of youth of color in this discourse, and the success that youth of color had in countering race-blind hegemonic narratives (captured through #AllLivesMatter) suggests that social media can serve as a democratizing force within the public sphere, which Habermas originally conceived of as exclusively accessible to the elite.

The Hashtag: #BlackLivesMatter as Counterpublic and Frame

Within the ongoing discursive struggles of the public sphere, #BlackLivesMatter simultaneously serves as a tool for generating a community of invested users and a framing strategy which is readily personalizable.

Hashtags are unique in the history of communicative innovations in that they combine the processes of classifying information and the developing social relations, creating what can be described as “searchable talk” (Zappavigna 2021). Hashtags allow users to identify and contribute to wider conversations, where the “conversation” does not resemble the turn-based mechanisms used in small groups but rather involves large groups of users entering in and out of a common narrative (Yang 2016). Even invoking a hashtag requires the presupposition of an audience to search for this hashtag (Marwick and boyd 2011). Given that the usage of a hashtag invokes an attachment to a community, it is unsurprising that different political groups invoke different hashtags; thus, usage of hashtags like #BackTheBlue and #PoliceBrutality is highly predictive of one’s political community online (Alfano et al. 2021). The observed formation of communities around hashtags aligns with the experiences of online Black Lives Matter activists, who describe the ability to connect with other activists and to amplify each other’s narratives as a major reason they use social media (Mundt, Ross, and Burnett 2018).

The community of #BlackLivesMatter hashtag adopters form a networked counterpublic with important discursive implications. One of the major functions of counterpublics connected to Black Lives Matter has been to question the politics of respectability, which discredits the experiences of those not deemed to have the appropriate self-presentation or fit “perfect victim” status (Hill 2018). For example, during the unrest in Baltimore after the death of Freddie Grey, users of the hashtag #BaltimoreUprising formed a temporary counterpublic which questioned narratives about urban unrest which played into racial fear mongering and ignored the harmful impact of police brutality on Black communities. This counterpublic was successful in forcing mainstream media outlets to alter their coverage (Welles and Jackson 2019). Increased mainstream media coverage is highly predictive of elite response to movement demands, so this discursive practice has offline significance (Freelon et al 2016).

Beyond generating community, #BlackLivesMatter also serves as an important framing strategy. When movement frames are used as hashtags, this provides the opportunity for "distributed framing" in which a large audience appends their own commentary to a movement frame, thus shaping the meaning of the frame (Ince, Rojas, and Davis 2017). For example, when users include #SayHerName in the same tweet as #BlackLivesMatter, they are highlighting the unique impact that police violence has on Black women, and this sentiment then becomes a part of the broader movement dialogue. Distributed framing is important in connective action, where the potential for development of personal identity is an important motivational force for engagement, because it provides a way for the individual to adapt movement messaging to their own experience. This is one way that digital technologies facilitate connective action. Individual users, though, are not the only ones capable of using distributed framing. Black Lives Matter organizations have used distributed framing to append additional frames to #BlackLivesMatter, with the most common addition focusing on individual rights, followed by Black culture, gender and income inequalities, and racial identity (Tillery Jr. 2019).

The specific properties of the hashtag also describe some of the limitations of hashtag activism. Black Lives Matter has been in usage for seven years, yet the majority of hashtag usage has occurred during peaks in the aftermath of widely shared incidents of police brutality. The visibility of trending hashtags on Twitter is engineered to fulfill users' desire for liveness --

“connection to shared social realities as they are happening” (Couldry 2004: 3) — as evident with features such as the “What’s Happening” trending section and the continuous news feed (Zulli 2020). Liveness induces a perceived “need to contribute to conversations as they happen”, which means that new external events may be frequently discussed in their aftermath but fade quickly in lieu of additional triggers (Zulli 2020). This may explain one ongoing puzzle in regard to Black Lives Matter: despite mobilizing a movement of unprecedented size in June 2020 (Buchanan, Bui, and Patel 2020), few policy concessions and primarily symbolic victories have characterized the aftermath of Black Lives Matter’s efforts (Taylor 2021). This apparent paradox has been noted in past digitally-fueled, decentralized movements, and can be understood through a framework of “capabilities” -- what protestors have the capacity to do -- and “signals” -- what signals those capabilities send to those in power (Tufekci 2014). When engagement with a movement is engineered to be episodic and transitive, a movement’s capabilities may be fundamentally limited.

In the preceding sections, we have done the following: first, we classified the Black Lives Matter movement according to Bennett and Sederberg’s connective/collective action framework as a hybrid movement. Then we articulated the role of the platform and the hashtag in this framework. Now that we have considered the specific limitations and possibilities afforded by the #BlackLivesMatter, we can pursue an analysis of the hashtag’s spread which recognizes the significance of diffusion. To do this, we will adopt a network-based approach.

Networks, Diffusion, and Contagion

Network Theory

The term ‘social network’ elicits “connotations of textiles, webs, and grids”, evoking a “powerful image of social reality” as composed of individuals connected via a dense network of invisible bonds (Scott 1988:109). Since the network metaphor crystallized into a sociological concept in the 1950s, scholars who study networks have argued for its centrality within the field of sociology (Scott 1988; Wellman and Berkowitz 1988). Granovetter powerfully argued in his (1973) study, “The Strength of Weak Ties”, that analyzing networks was invaluable for making the link between individual- and societal-scale phenomena; Granovetter’s classic study has remained at the center of the social network canon ever since (Lazer 2009). In sociology, the convention is to represent a network as a collection of entities (“nodes”) which most frequently but not exclusively represent individuals and which are linked by a set of relations (“edges”). Social network analysis represents a wide array of techniques but borrows heavily from mathematics’ graph theory due to the observation that many mathematical measures of networks, such as centrality, have sociological implications (Berkowitz 1988). There is some disagreement over whether network analysis represents a theory or a methodology, but to the extent that network analysis compels a form of reasoning defined by an emphasis on the “functions or properties of kinds of ties” combined with “topological reasoning” regarding social structure, we can classify theories as network-based (Borgatti et al. 2014:5). The mathematical roots of network analysis have positioned the discipline to take advantage of the recent explosion in both computing power and so-called “big-data”; thus, Lazer (2020) argues that we are undergoing a “second revolution” in network analysis defined by the rise of research conducted on networks of a previously inconceivable size.

The transmission of social phenomena across the edges in social networks has been described separately within several fields of sociology. The earliest empirical studies examined the diffusion of innovations and noted that individuals' adoption of such technologies as hybrid corn and prescription drugs depended on one's proximity to an earlier adopter (Ryan and Gross 1950, Coleman et al. 1957). The paradigm of diffusion can be traced back to the ideas of Gabriel Tarde, a nineteenth-century French sociologist (Rogers 1995). Tarde used the terms "invention" and "imitation" to describe the production of new social phenomena and their subsequent spread via interpersonal relationships. Tarde viewed these processes as central to social order, writing that 'Socially, everything is inventions and imitations' (Tarde 1890 [1903]:3, emphasis mine). Within social movement and collective action scholarship, there was a recognition beginning in the 1980s that whether one's peers have engaged in a behavior is a major factor in motivating and sustaining an individual engagement (Snow et al. 1980, McAdam 1986, Gould 1993).

Theorists observed, however, that the relationship between one's peers' participation in a behavior and an individual's behavioral adoption was often more complex than a linear function of the number of neighbor adopters. Again, Granovetter provides an early study to make this point. Granovetter (1978) observed that for socially risky collective behavior, an individual's propensity to join is mediated by how many of their peers have already joined. Each individual, Granovetter conjectured, has a unique "threshold" of participating peers which will trigger the individual adopting the behavior. In any population you will have instigators, willing to take risks even if none of their peers join, conservatives who will never join in on a behavior, and the rest of the population which has a threshold somewhere between these extremes. By running simulations with the threshold model, Granovetter demonstrated that minor variations in the distribution of individual thresholds could determine whether a collective behavior failed or saw widespread adoption in otherwise identical populations. The threshold model, though not explicitly a model of a network, reflects two critical ideas: first, individuals' decisions to participate in collective behavior are interdependent, and second, the use of mathematical formulas to model individuals' behavior can reveal otherwise-hidden insights on collective action dynamics. In models like the threshold model, these formulas reflect "formal and compact expressions of theoretical concepts" which can be used to develop theory and assess empirical findings (Borgatti et al. 2014:6).

Complex Contagion

It is vaguely evident that network analysis, theories of diffusion, and the concept of thresholds each have insights to offer our analysis of the spread of #BlackLivesMatter on Twitter, but in order to begin our analysis, we must identify a single framework which unifies these ideas. The framework we will adopt reflects Damon Centola's notion of simple and complex contagions. Centola (2007) describes a simple contagion as an entity (such as a behavior, belief, or use of innovation) which may be transmitted to susceptible individuals via exposure, where each exposure represents an independent probability of successful transmission. Centola recognized that those entities which are socially risky or which are only beneficial to adopters if they are popular, the simple contagion model fails, because it does not account for the threshold effect. Centola then defines complex contagions, which require "reinforcement from multiple sources" (Centola and Macy 2007:703) in order to flow between nodes.

The need for repeated activation from one's infected neighbors converts Granovetter's notion of a threshold into scalable theory, because where Granovetter's threshold model assumes that every individual is aware of the state of every other individual at any given time (impossible in any network of reasonable size), Centola's complex contagion model only requires an individual to be aware of the states of those individuals with which it shares a tie. This distinction furthermore allows complex contagions to be modeled upon social networks. Modeling contagion within a social network involves labeling nodes as either susceptible or infected, and occasionally a third category, recovered (Shakarian et al. 2015); for our purposes, we will focus on the basic susceptible-infected model. The simplest expression of the complex contagion function is a step function, where a given node's state is unactivated (susceptible) until the sum of its peers' activations reach the given threshold, at which point it becomes activated (infected). Simulating simple and complex contagions upon networks informs the analytical relevance of the distinction between types of social contagion. For example, the types of weak-tie heavy networks which excel at transmitting simple contagions fail dramatically at transmitting complex contagions, because weak-tie pairs are unlikely to share enough neighbors to reach the threshold of necessary exposures before transmission (Centola and Macy 2007).

Empirical Applications

Given the accessibility of online social network data, online social networks have provided an opportunity to test theoretical accounts of social contagion. Some evidence has accumulated that politically-oriented hashtags, like #BlackLivesMatter, behave as complex contagions. In the first large-scale analysis of hashtag contagion, Romero et al. (2011) examined the probability that a user would adopt a given hashtag after their k-th neighbor used the hashtag. Romero et al. (2011) found that political hashtags behaved like complex contagions, as "repeated exposures continu[ed] to have unusually large marginal effects on adoption", whereas other hashtags were either adopted upon one of the first exposures or never. Furthermore, political hashtags have been observed to exhibit "resonant salience", where events drive widespread adoption of a hashtag throughout the network, but a small group of committed users sustains hashtag usage in-between widespread diffusion (Barash and Kelly 2012). The small group of committed users may represent the critical mass which is necessary for complex contagions to go viral (Barash 2011). Resonant salience may exhibit itself in #BlackLivesMatter as the inner core of Black Lives Matter affiliated accounts which tweet in-between widely publicized incidents of police brutality (Freelon et al. 2016).

Not all studies agree, however, that political hashtags are complex contagions because they are political. Weng et al. (2013) argues that complex contagions are representative of most hashtags, which become structurally trapped within communities due to the need for social reinforcement, while viral hashtags go viral because they act like simple contagions. Similarly, Mønsted et al. (2017) found that hashtags behaved like complex contagions, after artificially spreading hashtags through a botnet and measuring whether the botnet's followers adopted the hashtags. That being said, the notion that all hashtag adoption represents a complex contagion lacks a strong theoretical explanation, and has been contradicted by other studies, like the study by Hodas and Lerman (2014) which found that all hashtags spread like simple contagions. This set of seemingly contradictory findings point to the challenges of observational contagion research, which must account for two interfering phenomena: opacity and homophily.

Opacity and Homophily

The "opacity problem" refers to the fact that individuals will not see every tweet issued in their follower network. While some authors (e.g., Romero et al. 2011) have operationalized exposure as any neighbor's tweet, and assumed that this assumption would not significantly affect results, multiple scholars (Lerman 2016, Berry et al. 2017) found that varying the percentage of tweets an individual is assumed to have read can lead to differing results regarding whether a contagion is simple or complex. Fink et al. (2016) developed a solution to the opacity problem by incorporating randomness in their model to simulate the chance a user saw a given tweet; however, this method has not yet been used to examine an online social movement.

Homophily refers to the widely accepted observation that people with similar characteristics tend to cluster together on a network (Khanam, Srivastava, and Mago 2022). Homophily can be broken down based on whether the homophily is among individuals sharing similar social positions (status homophily) or beliefs (value homophily); homophily can further be broken down by its cause, whether it is because the underlying pool of potential ties is homogenous (baseline homophily) or because one is predisposed to adopting ties with similar individuals (inbreeding homophily) (McPherson, Smith-Lovin, and Cook 2001). In the United States, race and ethnicity are among the potent characteristics in terms of homophily effects (Smith, McPherson, and Smith-Lovin 2014). In addition to race, political orientation has major homophily effects. Not surprisingly, homophily effects have been shown to affect the flow of political information on Twitter, causing both liberals and conservatives to be exposed to more like-minded information (Halberstam and Knight 2016). A related study found that there was "extremely limited connectivity" between right-leaning and left-leaning users in political retweet networks on Twitter (Conover et al. 2011). The cause of high baseline homophily for characteristics including race and political orientation can be ascribed to geography, given that "geography is the physical substrate on which homophily is built" (McPherson et al. 2001). Residential segregation among Black, white, and Americans of other races has been recognized as a feature of American geography; segregation based on class and political orientation have also been noted as recently increasing (Iceland, Weinberg, and Steinmetz 2002, Massey, Rothwell, and Domina 2009). Since an individual's social ties most often come from the communities and organizations they participate in, geography causes one's social ties to be even more homophilous than inbreeding homophily effects would predict.

Homophily is relevant to contagion because it provides an alternative explanation for contagion-like behavior. If specific clusters in a network become infected while others do not, this may be because these clusters contained a critical mass of infected individuals who activated their neighbors. Alternatively, this may be because these clusters contained individuals with similar characteristics that made them all more likely to adopt a behavior (Ogburn 2018; Shalizi and Thomas 2011). For example, De Choudhury et al. (2016) found that states with higher rates of police violence tended to have greater #BlackLivesMatter-related activity on Twitter, controlling for other factors. Given that Black Lives Matter is a movement for racial equality which has proposed major efforts to combat income inequality (Dennis and Dennis 2020), we can predict that race and income level would also impact one's participation in the movement, and cursory analysis of the demographics of #BlackLivesMatter participation has suggested that this is true (Olteanu et al. 2015).

Distinguishing between homophily and contagion represents an important methodological challenge. If support for #BlackLivesMatter appeared to spread across a social network as a contagion, but actually grew through adoption by similar individuals who would benefit from policing reform in a similar fashion, then both researchers and activists could make incorrect assumptions about support for the movement. Furthermore, overlooking individuals' salient characteristics could mean missing insights into how a contagion spread across different groups of individuals. For example, individuals in a wealthy, white neighborhood might be more hesitant to support #BlackLivesMatter, as they often benefit from high-levels of policing and racial inequality, whereas individuals in a neighborhood with a high rate of police violence could benefit from policing reform. Exploring homophily would mean identifying if some groups treated support for #BlackLivesMatter as a complex contagion while others treated support as a simple contagion, which would have major significance for organizers and our understanding of how individuals' decisions contribute to racial inequality.

Summarizing Findings

The literature demonstrates that online activism, rather than being an inconsequential and lesser-form of activism, reflects a modern way for individuals to engage in activism. Hashtags are powerful tools which shape movement framing, engender counterpublics, and allow individuals to debate policy and social meaning across a decentralized, low-commitment network of social media users. Hashtags' episodic nature may limit the influence of #BlackLivesMatter on lawmakers, but the hashtag still serves as a powerful discourse-shaping device. Centola's theory of simple and complex contagion may shed light on how social movement hashtags can spread across a network, but questions regarding homophily and opacity effects have limited past observational research of contagions on social media; thus, our analysis should make efforts to address these effects.

CHAPTER 3: RESEARCH QUESTIONS

In my research, I seek to understand #BlackLivesMatter's spread in the week after the death of George Floyd, a period of intense social unrest in the United States. I define three questions, but I do not articulate a hypothesis for my first research question, which is primarily intended as an exploratory question to shape the rest of my analysis.

RQ1: How did the #BlackLivesMatter hashtag diffuse on Twitter in the aftermath of George Floyd's death?

RQ2: How did hashtag adoption vary by neighborhood characteristics?

H2.1: Individuals from areas with high levels of population density, Black residents, racial income inequality, and violence will adopt #BlackLivesMatter more readily than other individuals.

RQ-3: How can social contagion explain the diffusion of #BlackLivesMatter?

H3.1: The #BlackLivesMatter hashtag is a complex contagion, requiring multiple exposures before adoption, but the contagion effect will be moderated when accounting for homophily among users.

H3.2: Users in more urban areas with higher rates of police violence and larger Black populations will require less social reinforcement to adopt #BlackLiveMatter han other users.

By answering these questions, we will be able to better understand how and why individuals join social movements, which has both practical implications (to activists seeking to improve tactics) and academic relevance (to sociologists seeking to understand movement growth).

In each of the following three chapters, I tackle one of these research questions. Because each of my research questions build upon each other, I define the methods for each question at the beginning of each chapter and then proceed to describe the results in the remainder of the chapter.

CHAPTER 4: METHODS AND RESULTS I: OVERALL TRENDS

To answer my first research question, regarding how #BlackLivesMatter diffused on Twitter, I describe my process of data collection and outline broad characteristics of my dataset. I examine how the characteristics of hashtag adopters evolved over time, finding that the diffusion event was seeded by a core of highly-active users. I also begin my geographic analysis by examining the global scope of #BlackLivesMatter and describe the implications of this. The findings of this chapter are critical for describing the overall diffusion event, and providing context for the following two chapters, when we focus on the narrower population of adopters who can be pinpointed to a geographic location in the United States.

Sharing Code for Reproducibility

In the following three chapters, I describe the steps I took to analyze my data in a broad sense. In order to see the specific steps which I took to retrieve, sort, and analyze my data, and in order to reproduce my methodology, I point individuals to the GitHub repository where my code is stored in full: https://github.com/ray-hc/blm_network_analysis

Dataset Description

I collected Twitter data for this project using the Twitter API, specifically using Twitter due to its reputation as a public sphere and its relevance to the Black Lives Matter movement (as described in the prior chapter) and its data availability. Twitter provides a specialized academic research API which allows researchers to download up to ten million tweets monthly. The Twitter API allows researchers to filter tweets based on posting date, text or hashtags within tweets, and whether a tweet contains a geotag, among other properties. Each queried tweet provides the author's user ID, which can then be used to query basic metrics about the user, such as their user bio, name, and handle, as well as their list of friends and followers. By contrast, the world's largest social network, Meta's Facebook, has seen comparatively little social contagion research because it has limited APIs available to researchers other than those within the company (Hatmaker 2021; State and Adamic 2015); Meta-owned Instagram similarly lacks APIs which provide network data (Meta, Inc. n.d.). The specific affordances of the Twitter API cause Twitter to be overrepresented in observational social movement research; separate research has confirmed that Instagram was also a powerful motivator of user engagement with #BlackLivesMatter during the study period (Chang, Richardson, and Ferrara 2021), and I leave further exploration of this to future studies.

I downloaded all Tweets containing "#BlackLivesMatter" (case-insensitive) in the week of George Floyd's death. This covered a time period beginning at midnight Eastern time on the day that George Floyd was murdered, May 25th, continuing until June 2nd. Because I mistakenly did not convert the query end-time into Coordinated Universal Time (UTC), this dataset does not include tweets sent between 8pm and midnight Eastern time on June 2nd.

For each user in my dataset, I downloaded the following information: birthdate, the number of users they follow, and their number of followers. Because the principles of ethical research state that we should collect only the data necessary for our analyses and because I was conducting a network-based analysis, where network structure is more important than individual characteristics, I did not store the actual content of tweets or the names of users. Additionally, some users had converted their accounts to private during the three-month gap between when I conducted my data collection for Tweets and when I downloaded user data, so I also checked account visibility. Any users who had converted to private were filtered out of my datasets when I began my analysis. For users with geotags, I additionally checked whether or not they previously used #BlackLivesMatter before the study period and download their Tweet rate per hour for the 30 days preceding June 2nd, the list of users they follow, and the embedded locations for all geotagged tweets they posted between May 25th and June 2nd.

During the study period, 14,176,614 public tweets containing #BlackLivesMatter were sent by 5,091,940 hashtag adopters, as shown in Figure 1 below. The number of cumulative adopters steadily rose over the study period, with the viral diffusion of #BlackLivesMatter beginning in the morning on May 26th, just over twelve hours after Floyd's murder. The initial peak in #BlackLivesMatter activity on May 28th coincided with the national spread of Black Lives Matter protests, and then more sustained #BlackLivesMatter activity occurred on May 31st and June 1st, a period which encompasses the infamous photoshoot by Donald Trump at St. John's Church, which was preceded by the violent clearing of protestors from the area (Taylor 2021).

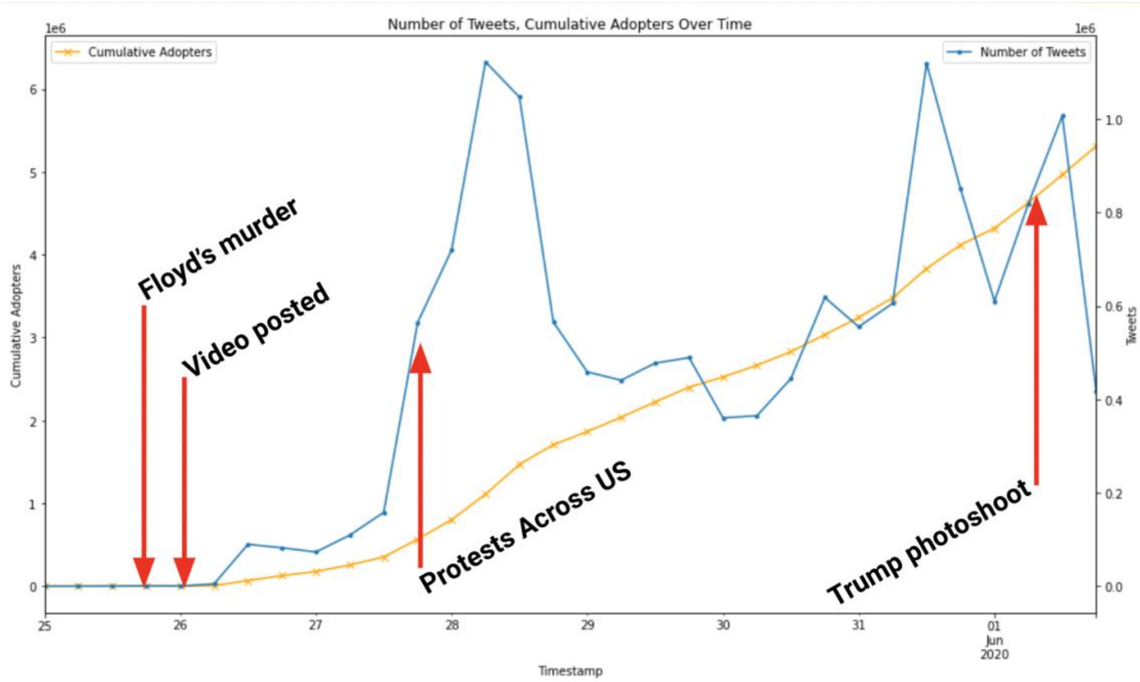

Figure 1: Cumulative adoptions and number of #BLACKLIVESMATTER tweets over time, annotated with important moments.

Comparing Early and Late Adopters

Next, I examine how user properties changed over the duration of the study period. The seed adopters of #BlackLivesMatter, those adopters who used the hashtag before the hashtag went viral midday on May 26th, tended to be more longstanding, highly-active, better-connected Twitter users than those who adopted thereafter. I use several metrics to gauge the behavior of Twitter users, specifically a user's number of followers, account birthdate (the date on which a user account was created), and number of Tweets. Account birthdate is insightful as older accounts may be more enmeshed in Twitter culture and have had more opportunities to be exposed to #BlackLivesMatter discourse, while number of Tweets (over the account's entire lifespan) indicates overall activity, and number of followers indicates directly how well-connected a user is.

At the start of May 26th, before significant adoption activity, the median birthdate of a hashtag adopter was January 2014, the median number of followers was 436, and the median number of lifetime tweets was 14,662. As the number of adopters began to increase rapidly, however, all measures of activity began to drop. By the start of May 28th, the median number of followers of a new adopter had dropped by nearly-half to 261 followers, the median number of tweets had dropped by one-third to 9,880, and the median account age had decreased by three years, with a median birthdate of January 2017. Given that the earliest adopters were much more active than the later adopters, I describe the earliest adopters as Twitter "power users".

In addition to being more active on Twitter, the earliest adopters tended to be preexisting supporters of Black Lives Matter. As visible in Figure 2, on May 25th, over 70% of adopters had previously used the hashtag; by June 2nd, less than 20% of new adopters had previously used the hashtag. This fact suggests that, not only were early adopters "power users", but they also previously supported the movement. These adopters, who we might consider to be the activist core that maintains #BlackLivesMatter between peak periods of attention, are crucial for the earliest phase of adoption.

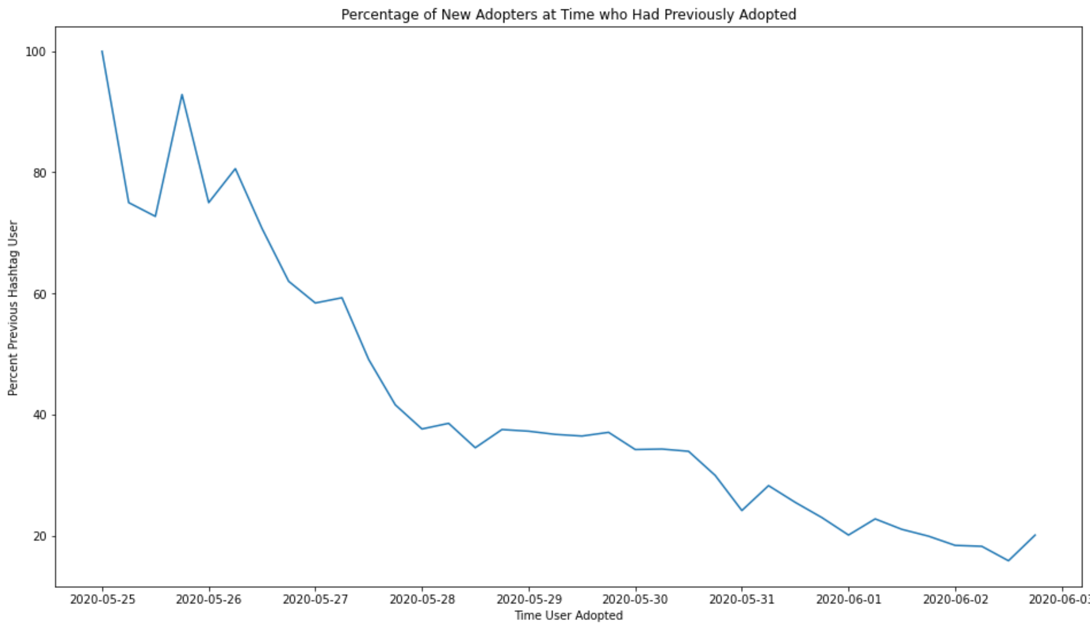

Figure 2: Percentage of those whose first usage within the study period of the hashtag occurred at a given time, who had previously used the hashtag (before the study period).

The behavior of the earliest adopters is important because since #BlackLivesMatter first went viral in response to the killing of Tamir Rice, its adoption has been characterized by periods of low engagement interrupted by sharp spikes in response to police killings and newsworthy events (Olteanu et al. 2015). Given that Twitter has been an important organizing ground since the inception of the Black Lives Matter movement, it seems likely that many of these highly-active, pre-surge users are activists or otherwise close to the movement. The fact that these seed adopters are so well-connected is important for any contagion processes, as a few well-connected nodes can quickly spread a contagion throughout a network.

The Global Scope of #BlackLivesMatter

The final analysis I conduct in order to answer my first research question is an examination of geotags by country. While I more thoroughly examine the geographic characteristics of #BlackLivesMatter in the next chapter, my initial analysis found that the #BlackLivesMatter diffusion event was a truly global phenomenon. It is uncertain how many non-geotagged tweets came outside of the United States, however, a narrow majority of geotagged tweets during the study period came from outside the United States. Of 2.6 million geotagged tweets, 1.36 million (51.4%) geotagged tweets came from outside the United States. The top foreign countries represented in geotagged tweets were Brazil (300,969 tweets), Great Britain (223,445 tweets), Canada (74,791 tweets), Nigeria (64,732 tweets), and South Africa (59,832 tweets). Many of the most well-represented countries had large Black populations, historical ties to the United States, and their own legacies of racial inequality. The large proportion of geotagged adopters from outside the United States suggests that foreign adopters had a significant impact on the overall adoption rate. These users tended to adopt later than the overall adopter population, as shown in Figure 3, suggesting that #BlackLivesMatter first spread domestically before reaching an international audience.

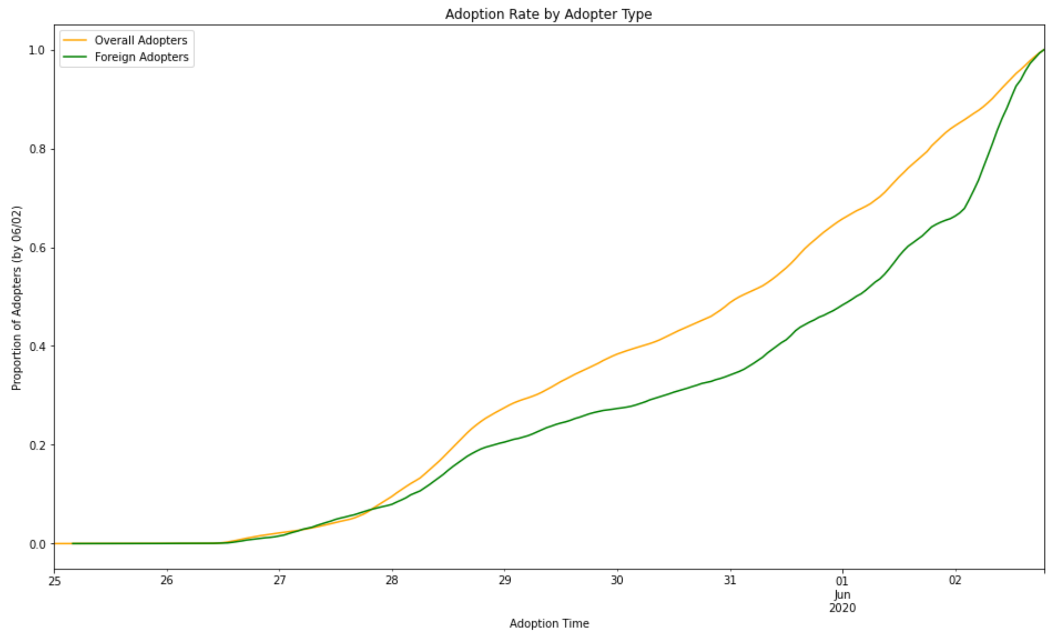

Figure 3: The overall cumulative adoption rate (as proportion of total adopters) for the overall adopter population, and for those adopters who have at least one tweet geotagged outside the United States.

Summarizing Findings

Now I answer my first research question:

RQ1: How did the #BlackLivesMatter hashtag diffuse on Twitter in the aftermath of George Floyd's death?

I find that hashtag adoption involved an unprecedented number of individuals and coincided with major offline developments regarding the movement. The diffusion event began rapidly, roughly twelve hours after George Floyd's death, and quickly rose in tandem with the spread of protests nationwide. The earliest period of hashtag adoption was driven by Twitter power-users, many of whom had previously supported #BlackLivesMatter, but as time went on, less-connected users new to the movement engaged with the hashtag. The diffusion of #BlackLivesMatter was a global event, including many individuals from outside the United States, and these international users tended to adopt later than domestic adopters. In this chapter, I have outlined only a birds-eye view of #BlackLivesMatter, but that is because the methods I describe in Chapters 5 and 6 allowed me to examine the behavior of a narrower population of adopters, those adopters whose geotags allow them to be geolocated within the United States, at much greater depth, in order to connect observed hashtag adoption dynamics to offline identities and to mechanisms of diffusion.

CHAPTER 5: METHODS AND RESULTS II: ADOPTION PATTERNS BY GEOGRAPHY

Next, I examine how adoption of #BlackLivesMatter varied based upon the salient characteristics of adopters' county subdivisions. In my literature review, I identified a need to examine the personal characteristics of hashtag adopters, because personal characteristics drive homophily, a confounding factor that needs to be considered in contagion research, and because individual identity often shapes engagement with social movements (e.g., Dixon and Roscigno 2003; Walgrave, Rucht, and Aelst 2010). Unfortunately, Twitter collects minimal information on demographic characteristics, which are one major source of homophily and critical for understanding an identity-oriented movement like Black Lives Matter. While Twitter lacks demographic information, it does provide another source of information: geotags. I use user geotags to examine geography because many of the properties which shape our social relations, such as race and income, also shape where we live (Iceland et al. 2002; Massey et al. 2009). Furthermore, the phenomena of racial inequality and police violence which has driven the Black Lives Matter movement's demands are heterogeneously distributed across the United States (Derickson 2017; Parks 2012). Therefore, I use the characteristics of geotagged adopters' county subdivisions in order to understand how salient geographic characteristics affect adoption. By examining both adoption rates and network properties, I find that among geolocated adopters, adoption was driven by users from urban areas, and that users from Minnesota and areas with large Black populations were important early adopters.

Using Geotags To Infer User Location

Past research has inferred Twitter user locations in order to study the spread of such diverse phenomena as forest fires (De Longueville, Smith, and Luraschi 2009) and misinformation campaigns (Jiang et al. 2020). Within my theoretical framework, I have predicted that race, income, and rates of police violence impact individuals' likelihood of supporting Black Lives Matter. Given that these characteristics are heterogeneously distributed across neighborhoods, we should be able to use the characteristics of adopters' neighborhoods in order to predict affiliation. Inferring neighborhood characteristics is not the only way to ascertain the impact of users' salient characteristics on user behavior, and I have included a longer discussion of alternative methods for determining salient characteristics in my supplementary materials on GitHub, but in summary, geotagging is a relatively easy and accurate method of stratifying users. In order to infer adopters' neighborhood characteristics, I examine the subset of adopters whose #BlackLivesMatter tweets are geotagged and use their geotags to infer the county subdivision in which they live.

Researchers interested in identifying users' locations based on Twitter data can choose between examining users' locations as defined in their bios, their tweets which contain general geotags (which represent places such as cities, states, and points-of-interest), and their tweets which contain precise geotags (which include GPS coordinates). User-defined locations are not an ideal choice for efficient geographic analysis because users can define any text as their location, whether or not it is a real place, and many users put fake locations in their bios (Hecht et al. 2011). As a result, I chose to infer user locations based on geotag.

For a user to have geotagged tweets, they must enable location sharing, either by selecting the option in settings or by selecting the geotag button when writing a new Tweet.

General geotags are the default option on Twitter and represent Twitter's best guess at which relevant location that a user is at, based on the user's GPS information. The process of choosing a general geotag to append to a tweet is shown in Figure 4. Twitter's locations are sourced from Foursquare, which maintains one of the largest geospatial datasets (Martineau 2019). Users are able to re-assign their geotag to an alternate location, as long as it is within the Foursquare database. Once geotagging is enabled, Twitter will automatically include a user's broad location in each subsequent tweet, until it is disabled, as shown in Figure 5.

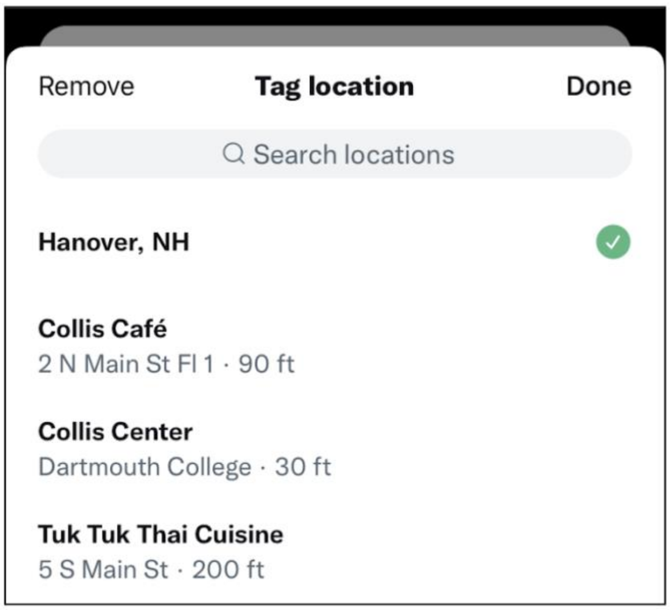

Figure 4: When I enable geotagging, my location defaults to Hanover. I may choose to select "Collis Café" or any other searchable Foursquare location.

Figure 5: In subsequent tweets, my current location will be automatically included, until I select the blue location icon and remove my location.

General geotags contain a bounding box of coordinates representing the geographic space enclosing the location. Precise geotags, on the other hand, contain the exact coordinates of a Tweet. In 2019, Twitter turned off the ability to manually add precise geotags to tweets. Precise geotagging can still occur when users share Tweets from third-party applications, or when users post photos using Twitter's camera (Hu and Wang 2020), but most geotags now use the broader locations provided by Foursquare (Benton 2019).

In order to geolocate #BlackLivesMatter adopters, I first downloaded the geotags for all geotagged #BlackLivesMatter tweets during the study period. These represented 0.60% of all #BlackLivesMatter tweets sent during this period. Roughly 10% of these collected geotagged tweets contained precise coordinates. In order to avoid reducing an already-small sample size, I chose not to specifically focus on precise geotags, even though precise geotags offer a high-level of granularity. Because Twitter automatically assigns the name and bounding box of a precise geotag's enclosing city to that geotag, some geotags which appear to be at the city-level in my data analysis may actually represent precise geotags.

Next, for each user with at least one #BlackLivesMatter geotagged tweet, I downloaded the locations from all of their geotagged tweets during the study period. Each geotag is assigned a place type, and the frequency of different place types in geotagged users' geotagged tweets in my study period is displayed in Table 1. Because the vast majority of geotagged tweets are at the city-level, I chose to group users by county subdivisions, in order to maximize both the number of geolocated users and the level of precision regarding their neighborhood attributes. As shown in Table 2, the topmost geotagged locations are all major US cities.

Type of Place	Number of Geotagged Tweets at Place Type
admin	259,109
city	2,220,881
country	99,443
neighborhood	3,619
poi	14,538

Table 1: Number of geotagged tweets by place type. This includes all geotags during study period by users with at least one geotagged #BlackLivesMatter tweet.

Name	Type	Country	Geotagged Users with Tweet at Location
Los Angeles, CA	city	US	3073
Manhattan, NY	city	US	1687
Brooklyn, NY	city	US	1174
Chicago, IL	city	US	1016
Houston, TX	city	US	875

Table 2: Top locations by number of users tagged at each location.

In order to convert geotags' bounding boxes into county subdivisions, I filtered out all geotags with locations of type "country" or "admin", as these do not provide precise enough geographic information to map users to county subdivisions. I filtered out all users with one or more geotagged tweets outside the United States, which represented roughly half of geotagged users.

Then I computed the centroids for all remaining users' geotags' bounding boxes, in order to compute a mean centroid for that user, weighting for the number of tweets per location. Then I computed the modal centroid for each users' geotags. If the mean and modal centroid are more than thirty kilometers apart, I removed the user from my population of geolocated adopters. For all remaining users, I computed a user's county subdivision as the county subdivision containing the user's modal centroid.

I believe this methodology to be an accurate method of identifying a user's home county subdivision for several reasons. First, the deviation between mean and modal centroid followed a power-law distribution, with the majority of users' deviation between mean and modal centroid not exceeding ten kilometers. Second, the median number of geotagged tweets per geotagged user was 11, so there are many data points with which to estimate location. Third, the study period captures eight days of geotagging information, and human mobility patterns follow a weekly rhythm (Hasan, Zhan, and Ukkusuri 2013). Finally, the study period was only three months into the COVID-19 pandemic, when many people were still in lockdown. Huang et al. (2020) examined the mobility of twenty-million users in the United States with geotagged tweets between March and May 2020 and found that user mobility remained significantly lower than the pre-March baseline. The one exception they found was in Minnesota during the last week of May, when many Black Lives Matter protests occurred in response to the murder of George Floyd. During this time, measures of mobility increased in Minnesota, but the average variation in user location from day-to-day remained below pre-March levels for all but one day during that week, suggesting that most travel during the study period remained local.

After completing the steps described above to pinpoint users to county-subdivisions, I had a population of 34,379 geolocated adopters.

Comparing Geolocated Adopters to Overall Adopters

In Appendix A, I examine how representative geolocated adopters are of the overall population of adopters in depth, but I summarize those findings here. Geolocated adopters tend to have older accounts and are better-connected on Twitter than other users. Geolocated adopters tended to adopt slightly earlier than other adopters. I only geolocated adopters inside the United States, so part of the reason that they adopted earlier than adopters overall can be explained by the fact that a greater proportion of later adopters came from outside the United States. Additionally, we identified in the prior chapter that the earliest phase of adoption was characterized by adoption by Twitter power-users; since geolocated adopters tend to have these characteristics (and are more "on" Twitter), they also demonstrate an earlier adoption rate. That being said, the geolocated cumulative adoption rate (as measured by the proportion of adopters who had adopted already, as compared to total adopters by June 2nd) never differed by more than twelve hours from the overall adoption rate. In summary, geolocated adopters are not representative of the overall population, but understanding their adoption patterns is still important, because their adoption rate does not differ drastically from other users, and because early adopters may have an outsize impact on later hashtag diffusion. Now that I have contextualized the representativeness of the geolocated adopter population, I begin my analysis of hashtag adoption by geography.

Predictive Characteristics for Hashtag Adoption and Network Centrality

In the first stage of my analysis for my second research request, I conducted a linear regression to identify the characteristics of subdivisions which were most predictive of hashtag adoption and user network centrality. To do so, I used the American Community Survey (ACS)'s data on income, racial makeup, and population density, as well as the Mapping Police Violence dataset (MPV)'s data on police killings (Campaign Zero 2022; Bureau 2020).

The ACS is a survey conducted by the United States Census Bureau to share statistics on households at multiple geographic levels of analysis. I used the Census' 2019 definition of county subdivision boundaries and their 2019 5-year statistical estimates for my analysis. From the ACS data, I retrieved the following features for each county subdivision: population density, the percentage and per capita income of individuals self-described as exclusively White, Black, or another racial makeup (including all individuals who described their race as Asian, Native, Pacific Islander, Multiracial, or some other race), and the percentage and per capita income of individuals self-described as Latino. Nearly all county subdivisions with geolocated adopters had all of these characteristics defined; however, some, like Black per capita income, were undefined if there was insufficient number of residents meeting that criteria in the county subdivision. Black per capita income was the feature with the most missing values, and 314 adopters lived in a county subdivision with no Black per capita income defined. For these missing values, the median value across county subdivisions was imputed for the purpose of clustering. Using the per capita income levels for Black and White populations, I additionally calculated a Black-White per capita income ratio, to gauge the level of racial income inequality in a county subdivision.

The MPV dataset includes the coordinates of all police killings since 2013, which I used to calculate the total number of police killings since 2013 in each county subdivision. I used this metric to compute police killings per square mile, and police killings per capita.

In order to determine the most salient characteristics for predicting adoption of #BlackLivesMatter, I conducted a linear regression to determine which characteristics were most predictive of high per capita adoption after standardizing all characteristics. I then used a machine-learning technique called Lasso regularization to progressively identify the least relevant characteristics for diffusion (Pedregosa et al. 2011). Lasso regularization is an alteration to the linear regression algorithm which applies a penalty for large coefficients. As the size of the scalar ("alpha") for the Lasso regularization's penalty increases, the least predictive feature's coefficient is reduced to 0, until only the most predictive feature is left.

Parameter	Coefficients
Percent White	-8.07E-06
Black-to-White Income Ratio	-2.70E-06
Percent Other Race or Multiracial	-1.91E-06
Overall Per-Capita Income	-9.88E-07
Percent Black	1.52E-06
Latino Per-Capita Income	2.04E-06
Percent Latino	3.37E-06
Black Per-Capita Income	3.39E-06
Police Killings per Square Mile	4.90E-06

White Per-Capita Income	7.31E-06
Population Density	1.70E-05

Table 3: The results of my regression before applying any penalty via Lasso regularization. Color-coded where absolute green is the most positive and absolute red is the most negative.

Parameter	Coefficients
Percent White	-3.86E-06
Black-to-White Income Ratio	
Percent Other Race or Multiracial	
Overall Per-Capita Income	
Percent Black	1.18E-07
Latino Per-Capita Income	
Percent Latino	
Black Per-Capita Income	
Police Killings per Square Mile	1.46E-06
White Per-Capita Income	1.57E-06
Population Density	1.43E-05

Table 4: The results of my regression with Lasso regularization with an alpha of 5e-06. Color-coded where absolute green is the most positive and absolute red is the most negative.

Using conducting regression and Lasso regularization, I found that population density is the most predictive feature for per capita adoption, followed by White and Black population percentage, White per capita income, and police killings per square mile. Table 3 illustrates the regression results before regularization while Table 4 illustrates regression with a moderate amount of regularization (alpha=5e-06). Upon increasing alpha to 1e-05, only population density continued to have a non-zero coefficient, making population density the most predictive characteristic of per capita adoption. I had originally used police killings per capita in my analysis, but Lasso regularization showed that this was not highly predictive of per capita adoption rates. Because a police killing is such a rare event at the level of county subdivisions, the data can be erratic. I hypothesized that police killings per square mile would be a more valuable metric because, while it may be unlikely for an average individual to know someone personally killed by police, any police killing in their area may make them more suspicious of police overall. An interesting detail revealed by Lasso regularization is that, despite Black Lives Matter's emphasis on income inequality, the income inequality ratio was not strongly predictive of the per capita adoption rate. Furthermore, the Black per capita income metric was weaker than the White per capita income metric, which I attribute to the increasingly large political divide between more-educated affluent white populations and working-class white populations in the United States (Zitner and DeBarros 2018). I also attribute the class divide among white populations to the general weakness of white identity, given its status as the invisible and "unexamined default racial category", which makes racial identity a poorer unifier among whites than among populations whose racial identity implies shared oppression (McDermott and Samson 2005:248). White per capita income may also reflect education levels and a more liberal area.

Because an individual's degree of centrality is relevant for contagion dynamics, I also conducted a linear regression to determine which features are predictive of a high number of followers. I found that per capita income was most predictive of a higher number of followers, with an increase by one standard deviation of per capita income predicting an increase of 141 followers. Black population percentage also increased the number of followers, while an increase in the Latino and non-White/Black population percentages decreased the predicted number of followers. It is feasible that areas with large Black populations are more connected on Twitter given past discourse about ‘Black Twitter’ (Brock 2012, Sharma 2013) and Black Americans’ higher self-reported usage of Twitter (Auxier and Anderson 2021). Additionally, it has been noted elsewhere that Twitter tends to attract an affluent user base, with high-income individuals much more likely to report having a Twitter account than low-income individuals (Auxier and Anderson 2021), so wealthier users may be more connected because more of their peers are on Twitter .

Given the results of my linear regressions, I identify population density, Black per capita income, White per capita income, Black population percentage, White population percentage, and police killings per square mile as the six most salient characteristics of county subdivisions influencing the diffusion of #BlackLivesMatter. With this information, I grouped county subdivisions using the statistical technique, K-Means clustering, in order to compare adoption-rates across types of county subdivisions.

Using K-Means Clustering

In K-Means clustering, given a desired number of clusters K, a dataset will be broken down into K clusters in hyperdimensional space, where each feature of the dataset represents a dimension of space. Upon initialization, a vector for each cluster will be randomly placed in this space, and each observation will be assigned to the cluster whose vector is nearest to it as measured by

Euclidean distance. During each successive iteration, the cluster vectors will be updated to decrease the mean distance between each observation and its cluster vector, and this will continue until successive iterations of the algorithm cease to significantly decrease this cost measure, known as "inertia" (Lloyd 1982).

When conducting the algorithm, I weighted each county subdivision by the total population of that county subdivision. SKLearn provides an easy-to-use implementation of the K-Means algorithm, which I used (Pedregosa et al. 2011). K-Means Clustering is non-deterministic, so successive iterations of the algorithm may lead to slight variations in the final arrangement of observations into clusters (scikit-learn developers n.d.). There is a trade-off in K-Means clustering between inertia and number of clusters; as you increase the number of clusters, your inertia rate will decrease, but too many clusters become uninterpretable. To conduct K-Means clustering on my dataset, I first ran the algorithm with all values of K between 1 and 20. My results can be found in Figure 6; and as shown in the figure, above K=6, increasing values of K yielded successively less improvements in inertia; thus, I chose 6 for the number of clusters.

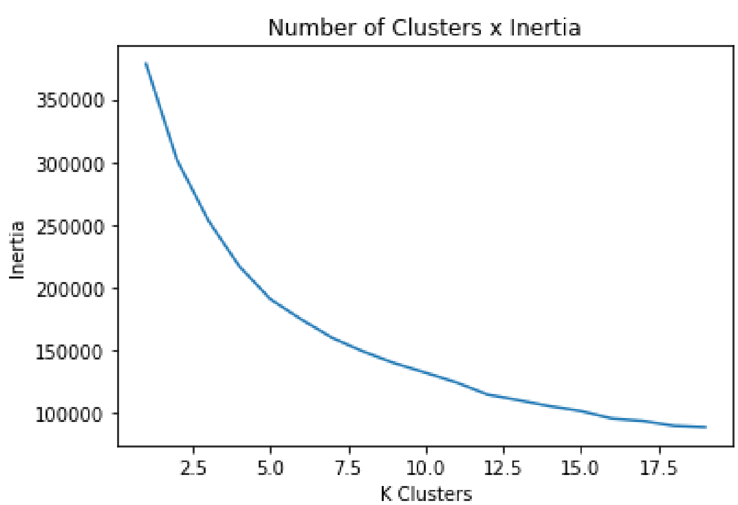

Figure 6: Examining the error rate of K-Means clustering based on K number of clusters.

After running the K-means algorithm to assemble six clusters of county subdivisions, I examined the characteristics of each cluster to determine the distinctive characteristics of each cluster, which I then named. First, I graphed the distribution of county subdivisions in projections of feature-space. One such projection can be found in Figure 7, where I examine each cluster by percent White, White per capita income, and population density. Then I attempted to find the most representative subdivision in each cluster by computing the weighted median for each feature and for each cluster, then calculating the Euclidean distance from each county subdivision within a cluster to the cluster's hypothetical perfect median subdivision. The county subdivision which came closest to the perfect median was deemed the most representative. Table 5 displays the most representative subdivision for each cluster.

After running these analyses, I named my six clusters based on what I deemed to be their most iconic characteristic(s): 1) the low-density, racially homogenous Rur.White cluster, 2) the medium-density Sub.Blacker cluster with its large Black population, 3) the medium-density, high-income Sub.Richer cluster, 4) the high-density BigCities cluster, 5) the SubRur.Poorer cluster, which spans the divide between the Rur.White and Sub.Richer cluster, and 6) the medium-density MidCities cluster, which is most notable for encompassing Minneapolis and the surrounding area.

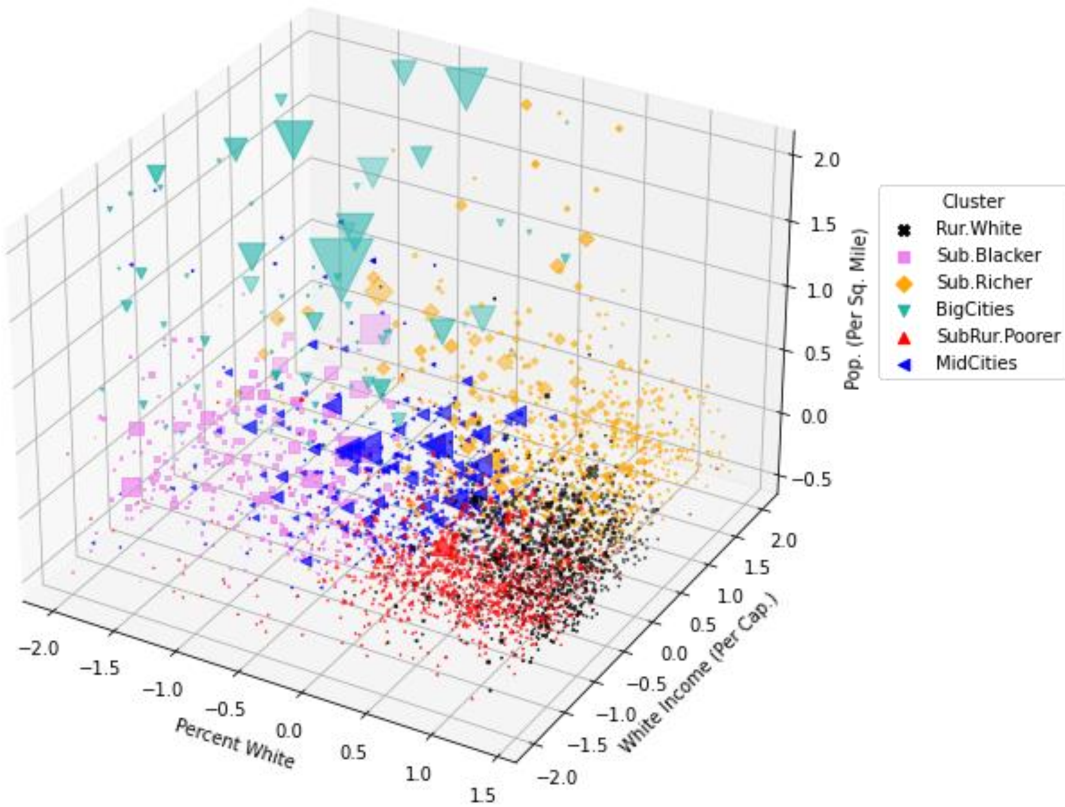

Figure 7: Each county subdivision with at least 1 adopter, plotted as a point in a three-dimensional space, where the axes represent percent White, White per capita income, and population density. The size of each point is proportional to the square root of the number of adopters from that county subdivision.

Most Representative Subdivision per Cluster						
Cluster	Name	Per-Cap. Income	White Pop. Percent	Black Pop. Percent	Pop. Density (per sq. mile)	Police Killings (per sq. mile)
Rur. White	Petersburg, North Dakota	$31k	91.2%	1.8%	165	0.00

Sub. Blacker	Airport, Missouri	$25k	50.5%	40.9%	2197	0.06
Sub. Richer	Burlington, Massachusetts	$51k	74.3%	4.5%	2356	0.00
BigCities	Chicago, Illinois	$37k	50.0%	29.6%	12,021	0.37
SubRur. Poorer	Alpine, Michigan	$28k	85.4%	6.3%	388	0.00
MidCities	Maplewood, Minnesota	$34k	68.2%	10%	2388	0.06

Table 5: Each cluster's most representative subdivision, as defined above.

Adoption Rates by Geography

After identifying my clusters, I examined how adoption rates and network properties by cluster varied by cluster. Overall, I find support for the hypothesis that more urban, diverse regions tended to be the strongest adopters of #BlackLivesMatter on Twitter, with individuals from big cities sharing #BlackLivesMatter on Twitter at far higher rates than individuals from other regions. Furthermore, I find that adopters from Minnesota, where George Floyd's murder occurred, were an important early source of adoption, as were adopters in areas with large Black populations.

First, I examined adoption rates by region, and I found that all regions in the United States had relatively similar adoption rates, with the exception of the Midwest, due to the early rise in tweets from Minnesota. As seen in Figures 8 and 9, users from across the United States adopted at relatively similar rates across regions, but users from Minnesota (part of the West North Central census division) adopted much earlier. By the beginning of May 28th, nearly 27%

of the Minnesotan adopters had already used #BlackLivesMatter once in the period, but the same was true of just under 12% of all geolocated adopters. Thus, Minnesotans were an important early driver of hashtag adoption.

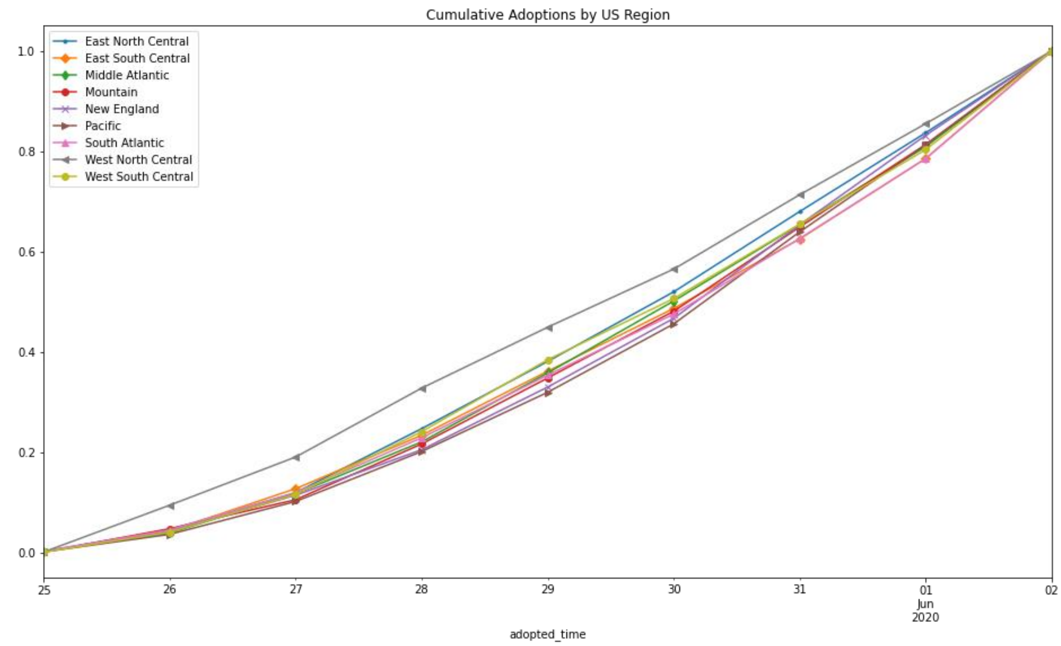

Figure 8: Cumulative proportional adoption rate by US Census region.

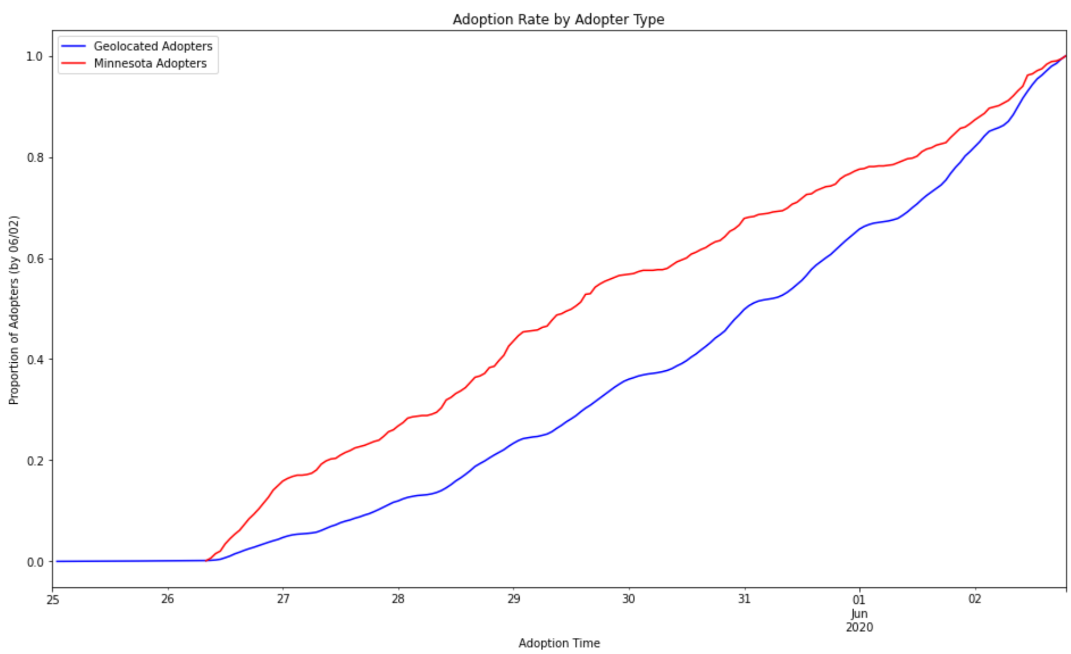

Figure 9: Cumulative proportional adoption rate for Minnesotan adopters compared to all geolocated adopters.

Next, I analyzed my previously-defined clusters, and found that salient subdivision characteristics have a pronounced impact on per capita adoption rates, with less White and more urban clusters adopting at much higher rates. I found that individuals living in BigCities county subdivisions were 9.84 times more likely to adopt #BlackLivesMatter in the study period than individuals living in Rur.White county subdivisions. While users from Minnesota were important early adopters, BigCities adopters dominated the discourse throughout the study period, as shown in Figure 10. After BigCities, the next highest adopting clusters were the Sub.Blacker cluster and the MidCities cluster.

It is worth noting that the per capita adoption rate captures the probability that an individual from a given cluster forms a Twitter account, enables geotagging, and uses #BlackLivesMatter. While it is not possible to determine per capita Twitter usage by cluster with my dataset, a Pew Research survey found that self-reported urban residents were 50% more likely than self-reported rural residents to have a Twitter account (Auxier and Anderson 2021). The same survey found that Black individuals are roughly 30% more likely than White individuals to use Twitter. Thus, some but not all of the higher engagement of more urban and diverse clusters can be accounted for by variable Twitter usage.

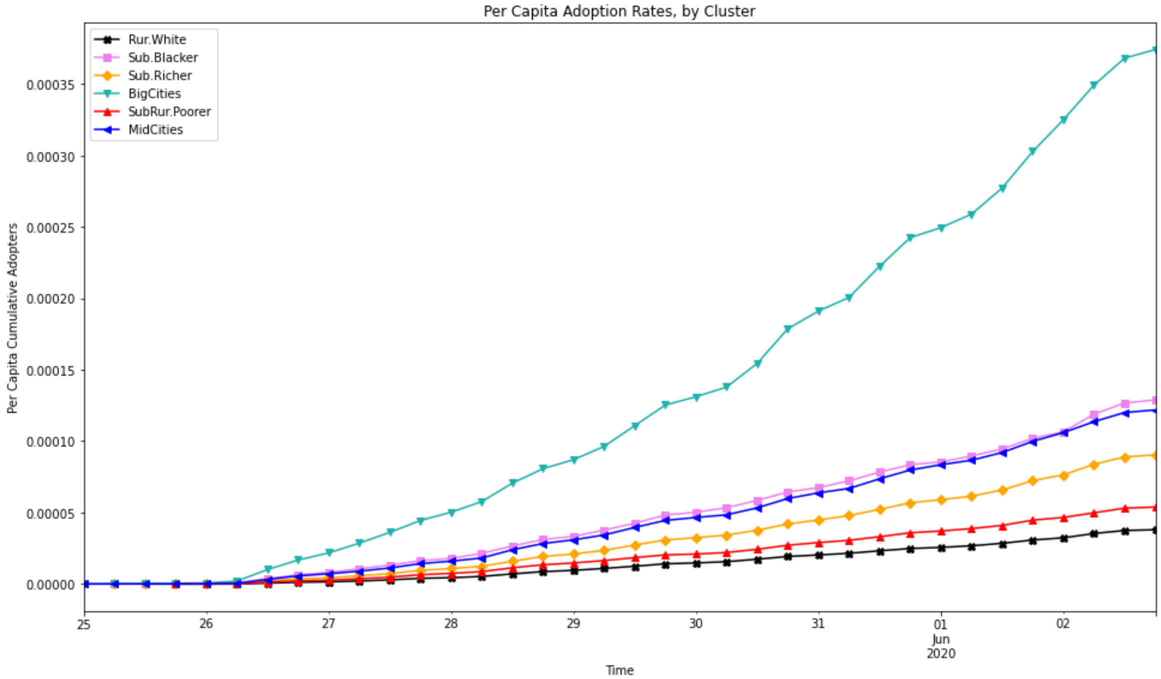

Figure 10: Per capita cumulative adoption rate by cluster.

Next, I normalized the adoption rates for each cluster based on the final number of adopters at the end of the study period in order to identify which clusters' adopters adopted earlier. I find that there is a minor difference between proportional adoption rates over time across clusters, with more urban and diverse clusters adopting earlier. When examining the full study period, I was surprised to find that the proportional adoption rates for each cluster were closely aligned. The divergence in proportional adoption rates was always limited to within six hours, which is less than the maximum variation between geolocated adopters and the overall adopter population. Focusing on the first 36 hours of major hashtag spread does reveal that Sub.Blacker and MidCities adopters adopted early at the highest rates, as shown in Figure 11. It is not surprising that the MidCities cluster adopted earlier, given that it contains Minneapolis. The fact that the Sub.Blacker cluster had the highest proportion of early adopters is notable, given that after previous instances of police brutality, "Black Twitter" was crucial for directing national attention to the issue (Hill 2018). The least diverse clusters, Sub.Richer, SubRur.Poorer, and Rur.White, saw lower rates of proportional adoption until near the end of the day on May 27th, when SubRur.Poorer's proportion surpassed that of MidCities and BigCities. After May 27th, there was no obvious pattern in the order of proportional adoptions. The earliest period of adoption is especially important because by the end of May 27th, #BlackLivesMatter had saturated the network; over 80% of new adopters thereafter followed at least four people who had already adopted within the study period (this finding is discussed in greater depth in Appendix B). In summary, adopters from areas with large Black populations were important early adopters, and the early adopters from MidCities and BigCities subdivisions serve as further proof of the importance of individuals from urban areas and from Minneapolis in explaining #BlackLivesMatter.

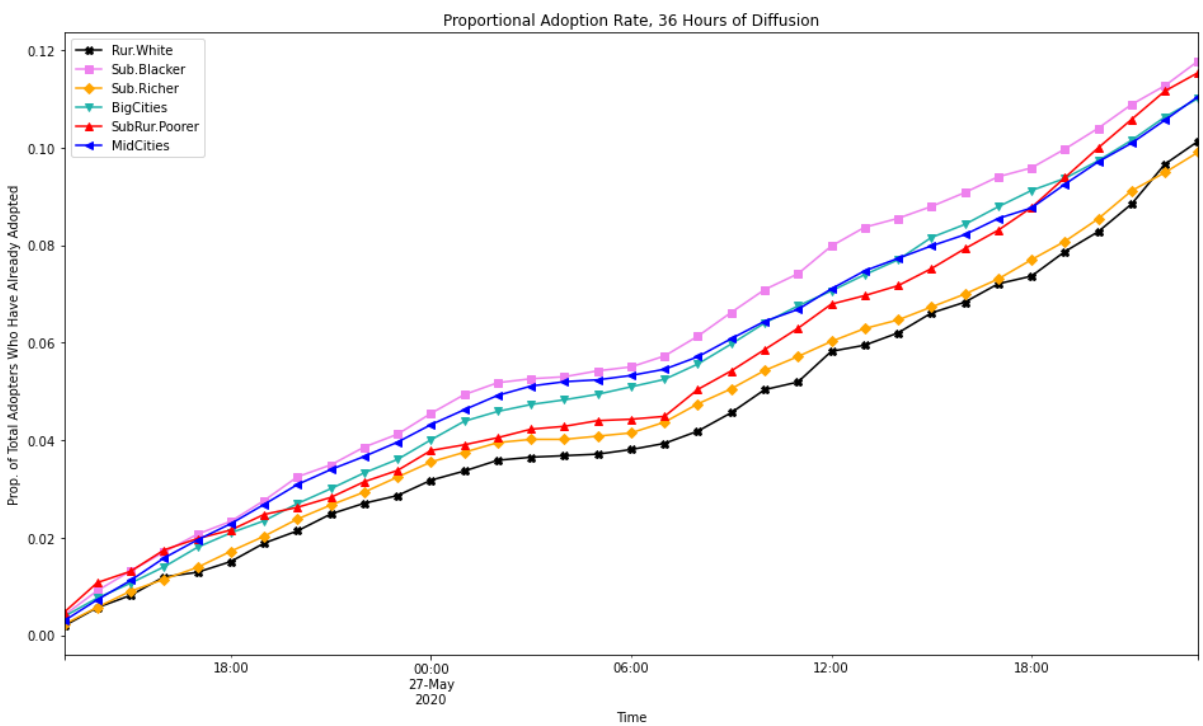

Figure 11: Cumulative Proportional Adoption Rates during first 36 hours of diffusion event, a time period covering noon on May 26th to midnight on May 27th.

Network Properties by Geography

For the final stage of my analysis for my second research question, I examined how network properties vary by geography, and I found that users from the BigCities cluster dominate the network of #BlackLivesMatter adopters. Furthermore, many of the most connected accounts came from BigCities clusters.

Examining who is followed by each cluster can illustrate where each cluster of users gets their social reinforcement. Table 6 illustrates where each "target" cluster received its information on Twitter, by measuring the proportion of geolocated adopters they followed that came from each "source" cluster. Every cluster followed adopters in the BigCities cluster more than adopters from any other cluster. As evidenced in Table 6, BigCities adopters' voices dominated other adopters' Twitter feeds. It is worth noting that the magnitude of this effect is partially due to BigCities adopters representing the largest cluster of adopters. An alternative metric is tie-ratio, which examines the number of ties between each cluster compared to the overall number of ties they could share (i.e. if every person in the target cluster followed every person in the source cluster). As is described in greater depth in Appendix B, an analysis of tie-ratios demonstrates that every cluster maintained the densest quantity of following-relationships with itself and the BigCities cluster, a finding which is not obvious in Table 6. The BigCities cluster was also the densest in-terms of the highest within-cluster tie-ratio, while the Sub.Blacker cluster was the second-most dense. The implication of these findings is that the BigCities had not only the highest per capita adoption rate, but BigCities adopters were also the dominant source of peer influence for other adopters.

Source Cluster:	Rur.White	Sub.Blacker	Sub.Richer	BigCities	SubRur.Poorer	MidCities
Target Cluster:						
Rur.White	0.10	0.11	0.13	0.38	0.08	0.21
Sub.Blacker	0.06	0.26	0.09	0.34	0.07	0.18
Sub.Richer	0.07	0.10	0.16	0.43	0.06	0.19
BigCities	0.04	0.07	0.10	0.60	0.04	0.14
SubRur.Poorer	0.08	0.12	0.11	0.36	0.13	0.21
MidCities	0.06	0.10	0.11	0.36	0.07	0.31

Table 6: A table illustrating, for each “target” cluster, the proportion of geolocated adopters they followed that came from each "source” cluster.

To further understand these network dynamics, I examined the median number of followers per cluster, and found that the BigCities cluster had more highly-influential accounts than other clusters. While the median number of followers for each cluster is relatively similiar (ranging from 373.5 followers for Rur.White adopters to 544 followers for BigCities adopters), the mean number of followers for a BigCities adopter is 5884.2, more than double the next highest mean follower count (2598.8 followers for Sub.Richer users), which suggests that a few extremely well-connected accounts help drove the dominance of the BigCities cluster. That being said, BigCities users still had the highest median followers, suggesting that these users are more well-connected overall.

In addition, I examined the top adopter accounts followed by each cluster, and found that each cluster tended to follow the same top accounts. I did so by identifying the top five most-followed geolocated adopters and top five most-followed adopters overall by cluster. The top adopters followed by each cluster were prominent Democratic politicians while the top geolocated adopters were a more eclectic group of liberal and lifestyle influencers. While the ordering of the top five varied, each cluster's top-five list included most of the same accounts. The only cluster which served as an exception to this statement was the Sub.Blacker cluster, whose top geolocated accounts included a few prominent Black activists. The rank order of top accounts did vary, but the top-five lists were almost identical by cluster. To limit ethical concerns, these lists are not included within this paper but can easily be reproduced by researchers who reproduce my data collection.

When examined together, my findings regarding the network properties of the #BlackLives have several implications. The dominance of the BigCities cluster along with similar Twitter diets across clusters can help explain why each cluster had a similar proportional adoption rate, because many of these adopters were receiving similar information from the accounts they followed on Twitter. The similarity of information sources also suggests that there is a degree of unmeasured homophily across clusters; in other words, geolocated adopters tend to be similar as a group. Finally, the dominance of the BigCities cluster and the relative strength of the Sub.Blacker and MidCities clusters, suggests that online engagement with Black Lives Matter is a predominantly urban phenomena, with adopters from areas with large Black populations also functioning as important contributors.

Summarizing Findings

Now I answer my second research question and assess my hypothesis:

RQ2: How did hashtag adoption vary by neighborhood characteristics?

H2.1: Individuals from areas with high levels of population density, Black residents, racial income inequality, and violence will adopt #BlackLivesMatter more readily than other individuals.

My hypothesis was partially true. Overall, I found that my specific metrics of racial inequality were less salient than broad demographic characteristics for determining hashtag adoption. Racial income inequality did not have a significant impact on hashtag adoption, and police violence (concentrated geographically in cities) had a slight impact but this was less salient than population density itself. This finding does not mean, however, that racial inequality is not relevant to the spread of Black Lives Matter: the map of modern America has been deeply shaped by racial inequality (Derickson 2017), a fact that I discuss in greater depth in my discussion. The most important finding from this chapter is that online Black Lives Matter activity is driven by urban areas, and that users from large cities dominate the online Black Lives Matter discussion. Without the involvement of activists and citizen journalists reporting evidence of George Floyd's murder, it may never have become a national movement, at least not quickly, but once news of Floyd's death spread, it was adopters from urban areas who drove the saturation of Twitter with movement messaging. Still left unanswered is why these adopters engaged with #BlackLivesMatter at the scale they did — was a process of online contagion facilitating the diffusion of #BlackLivesMatter, or can factors external to Twitter expalin hashtag adoption better? How did these mechanisms vary by cluster? Answering these questions is the goal of my next and final results chapter.

CHAPTER 6: METHODS AND RESULTS III: CONTAGION SIMULATIONS

Now I seek to apply theories of social contagion to the observed diffusion of #BlackLivesMatter. To do so, I use previously-defined methodology with some slight modifications. In two papers published in 2015 and 2016, Fink et al. examined evidence for complex and simple contagions, first by constructing four metrics with which to observe contagion in hashtag diffusion events (2015) and then by designing simulations to analyze whether hashtag diffusion events represented simple or complex contagions (2016). Though I do calculate Fink et al.'s four metrics for identifying contagion (2015), in general I found them to be a poor method for determining whether #BlackLivesMatter represents a simple or complex contagion. Their results were muddled, and my simulation results largely contradicted any evidence of contagion that they provided, but I still include these results in Appendix B for thoroughness and because these metrics also illuminate additional characteristics of the network of geolocated #BlackLivesMatter adopters. In this chapter, I define and then conduct a series of simulations which suggest that, over the entire study period, hashtag diffusion can be better explained by homophily and factors external to Twitter than by contagion processes; however, the first forty-eight hours of adoption indicate #BlackLivesMatter usage being driven by online contagion, which may imply a reciprocal relationship between Twitter diffusion and external news coverage.

Model Design

In this section, I define the models I used for my simulations. While I outline alternative model designs in my supplemental materials on GitHub, the models I define for my analysis are primarily based upon the models constructed by Fink et. al. (2016), with the primary variation being that I construct an additional component to represent the effect of homophily — whereby users with a similar propensity to adopt a behavior do so on the basis of an external influence.

Fink et. al. (2016) do not factor in the potential for external factors to influence hashtag spread. They can do this primarily because they only examine the first twenty-four hours of a hashtag's lifetime; thus, they assume that most users' adoption was due to online influence. There are two reasons why it would be valuable to expand upon their definition, however. First, even hashtags in the first twenty-four hours of diffusion can be shaped by offline influence; an example is the hashtag #MH370 which they examined and which was undoubtedly influenced by news reports about the missing Malaysian plane. Second, the social definition of hashtags can evolve over time; as users engage with #BlackLivesMatter, they shape its meaning (Ince et al. 2017). As a result of this evolution, examining the first twenty-four hours of #BlackLivesMatter may not provide much insight into #BlackLivesMatter today.

In order to account for external factors, I introduce an additional term in the contagion probability equations, which I call *externalProb*. I also introduce the usage of a reference model to handle the case in which there is no contagion at all – where hashtag adoption is determined entirely by external factors rather than online influence. Thus, in the following sections, I will define three models:

- The Homophily ("Reference") Model simulates the expected adoptions in the absence of online contagion effects, with adoptions motivated by volume of phrase usage in the public sphere and a group of users' propensities to adopt.
- The Homophily Plus Simple Contagion ("Simple") Model simulates users adopting due to a combination of external factors and simple contagion on Twitter.
- The Homophily Plus Complex Contagion ("Complex") Model simulates users adopting due to a combination of external factors and complex contagion on Twitter.

But first, I must explain how a core component of the contagion models, the random login schedules, are computed.

Constructing Log-In Schedules

Both the "Simple" and "Complex" models require the generation of login schedules for users. Randomized login schedules are important because they address the opacity issue in past contagion research, whereby assuming users are always on Twitter and getting exposed to hashtags in realtime can change whether a phenomena appears to be a complex or simple contagion (Berry et al. 2017). By testing models using multiple random login schedules, I can determine the reliability of my findings and more accurately reflect user behavior. Fink et al. (2016) used a Poisson model to generate random login times for hashtag adopters. Their Poisson model was based on the population-wide average number of logins per time-step, where a login was operationalized as a user tweeting at least once during the time period and time steps were 15-minute intervals. Fink et al. (2016) note, however, that Twitter activity levels follow a power-law distribution; to avoid any unnecessary assumptions, I use a user's actual average activity rate to generate random logins per user. For simplicity, I used 1-hour instead of 15-minute time-steps because this is the default granularity that the Twitter API provides when accessing a user's tweet counts. Additionally, to avoid the risk of external influence artificially inflating users' activity rate during the study period (e.g., in a scenario in which users are being motivated to tweet by offline events, like protests), I calculated the average login rate based on a user's previous 30 days of activity, ending on June 2nd. I found that the median login-rate was 0.058 logins per hour, which translates to 1.4 logins per day.

The Homophily ("Reference") Model

A reference model is the modeling-equivalent of a null hypothesis (Hobson et al. 2021). Its purpose is to determine if we can adequately explain the dynamics of #BlackLivesMatter adoption without the presence of online contagion. The reference model simulates the adoption which we would expect to occur even in the absence of contagion on Twitter.

If users are not motivated to tweet #BlackLivesMatter due to contagion on Twitter, it is likely that they are motivated by information obtained from the media or from their offline peers.

I use the rate of news articles containing "BlackLivesMatter" (as well as the variations, "blacklivesmatter" and "Black Lives Matter") to estimate the relative daily intensity of #BlackLivesMatter dialogue in the public sphere. I denote the news rate as d where d represents the number of days after George Floyd's death. I use Nexis Uni, a research tool for finding news articles, to determine the news rate (Nexis Uni n.d.).

Next, we define a parameter q_i which represents the propensity of users in a population to use the hashtag. Each user has their own set of underlying beliefs and behaviors which affects their likelihood of engaging with #BlackLivesMatter, given a certain level of #BlackLivesMatter discourse in the public sphere, operationalized as d. q_i represents the population's average propensity for adoption and may vary depending on whether the examined population is a specific cluster or the total population of geolocated adopters.

Using these parameters, I construct the function $externalProb$ to represent the probability of adoption due to external (not Twitter contagion) factors in a given hour on day d as:

$$externalProb(d) = \lambda_d \cdot q_i$$

Because this is the reference model, I assume external factors are the only factors affecting adoption. Thus, the overall probability of adoption in a given hour of day d is:

$$P(d) = externalProb(d)$$

This definition of $P(d)$ allows us to incorporate both the overall visibility of #BlackLivesMatter and group-level variation in attitudes towards the movement into a model of hashtag adoption, but crucially, it does not consider information about an individual's exposures.

To implement this model, I use Bernoulli trials. A Bernoulli trial refers to a random experiment with exactly two-outcomes, with a set probability of each outcome occurring. For each hour, if a user who has not previously adopted is deemed to be logged in, I run a Bernoulli trial to determine if there is a new adoption, where $P(d)$ represents the probability of adoption.

The Homophily Plus Simple Contagion ("Simple") Model

Following the approach of Fink et al. (2016), I model the probability of a user adopting due to contagion when they have been exposed to the hashtag k times as:

$$simpleProb(k) = 1 - (1-p)^k$$

where p represents the probability of adopting after a single exposure, and $(1-p)$ represents the probability of not being infected after an exposure. $(1-p)^k$ represents the likelihood of not being infected after k exposures, and $simpleProb(k)$ represents the opposite: the likelihood of being infected after k exposures.

The probability of a susceptible user adopting the contagious behavior during a given hour in which they have a login is defined as:

$$P(k,d) = simpleProb(k) + s \cdot externalProb(d)$$

I assume in the "Simple" model that both external factors and simple contagion influence users' ability to post; thus, s represents the proportion of first adoptions which are still caused by external factors when accounting for simple contagion effects.

To implement this model, for each hour, if a user is deemed to be logged in, I calculate the number of exposures since last login, then run a Bernoulli trial to determine if there is a new adoption.

The Homophily Plus Complex Contagion ("Complex") Model

The major distinction between simple and complex contagion models is that, in a complex contagion model, the probability of adoption after exposure is not fixed. Instead, as the number of adopters followed by an individual increases, that individual's probability of adoption after an exposure increases. We can use a logistic sigmoid function to model the varying probability of a complex contagion. The most basic version of this function is:

$$p_k = \frac{1}{1 + e^{-k}}$$

This version of p_k ranges from 0 to 1. A generalized form which can fit any shape and threshold is:

$$p_k = p_{lo} + \frac{p_{hi} - p_{lo}}{1 + e^{-g(k-k_o)}}$$

where p_{lo} and p_{hi} represent the minimum and maximum odds of adoption after a single exposure, g controls how steep the sigmoid function is, and k_0 controls the threshold of the function.

Graphically, this statement looks like this:

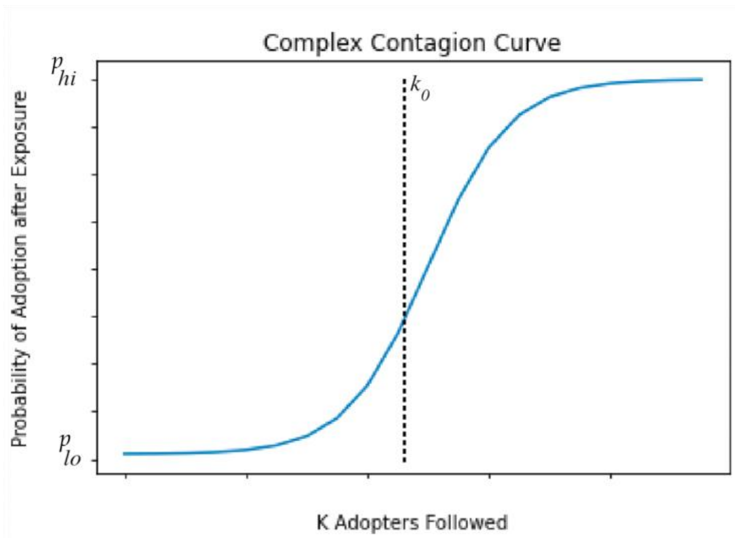

In order to connect p_k into our adoption probability function, some additional notation is required. Consider a list of exposures L for user u:

$$L_u = (e_a, e_b, \ldots, e_n)$$

Where the value e_j at each list position j represents the number of exposures to #BlackLivesMatter when the user had j adopting neighbors. The probability of adopting at some point in the sequence L_u due to complex contagion is defined as:

$$complexProb(L_u) = 1 - \left[\prod_{k=1}^{|L_u|} (1 - p_k)^{L_u[k]}\right]$$

Where p_k is the probability of adoption after a single exposure when a user has k adopting neighbors, and $L_u[k]$ represents the number of exposures a user had when they had k adopting neighbors.

Thus, the probability of adoption in an hour in which the user has a login is:

$$P(L_u, d) = complexProb(L_u) + s \cdot externalProb(d)$$

As in the simple contagion model, s represents the proportion of adoptions due to outside factors.

To implement this model, for each hour, if the user u is deemed to be logged in, calculate the number of exposures since last login, and reconstruct the sequence L_u in order to run a series of Bernoulli trials to determine if there is a new adoption.

Model Validation and Tuning

In order to run my diffusion models, I first had to compute the ideal parameters for each model. To do so, I adopted a simple grid-search, where I test all possible combinations of a predefined set of parameter options. Model performance is tested using the least-squares cost function, which is commonly used in machine learning and linear regression functions. For each hour in the time period, the difference between the predicted and actual number of cumulative adopters is squared, and the average squared error is the least-squares cost.

Each additional parameter to test for increases the number of options exponentially; thus, I was limited in the fine-tuning of models. In order to determine the set of options to test for each parameter, I first tested combinations of model parameters manually to identify rough upper- and lower- bounds for each parameter. Then, for each parameter, I refined the parameter options to a list of evenly-spaced values between the parameter's value for the best and second-best performing combination of parameters. The parameter options tested can be found in Table 7.

For the simple and complex models, I fixed several parameters to reduce the duration of the search process. For the reference model, there is a single parameter to test, q, which reflects the population's propensity to adopt given the news rate. Once I identified the optimal q in the reference model, I re-used this value in the simple and complex models. For the simple model, the additional parameters to search for are p (probability of adopting after single exposure) and s (proportion of adoptions caused by external factors). The probability of adoption after an exposure in the complex model is calculated using the sigmoid function, which requires the parameters, ϵ_{lo} (minimum probability of adoption after exposure), p_{hi} (maximum probability), k_0 (threshold), and g (how quickly p changes). Fink et al. (2016b) fix p_{lo} at 0.001, p_{hi} at 0.25, and g at 1. I re-use their value for p_{lo} at 0.001, but I include p_{hi} among my parameters to search for. I set g to 0.15 based upon early trial-and-error in which I found values between 0.1 and 1 tended to be a better fit. When g=0.15, the sigmoid function's rise from p_{lo} to p_{hi} occurs over a range of roughly 50 adopters, roughly 10% of the median number of friends for a geolocated adopter. This range is wide enough to allow for variations in individual adoption thresholds while maintaining the threshold shape. Even when fixing the values of q, p_{lo}, and g for the complex model, there are still three parameters to search for: p_{hi}, k_0, and s. For each model, I identified six options for the primary parameter (q for the reference model, p for the simple model, and $k_0$for the complex model) and five options for all other varying parameters.

Reference Model Parameter Options	
Round 1	q: [0.000001, 0.000005, 0.00001, 0.000015, 0.00002, 0.00005]
Round 2	q: [0.00001, 0.000011, 0.000012, 0.000013, 0.000014, 0.000015]
Simple Model Parameter Options	Fixed: q=0.000013
Round 1	p: [0.0001, 0.001, 0.003, 0.005, 0.01, 0.015], s: [0.1, 0.3, 0.5, 0.7, 0.9],
Round 2	p: [0.001, 0.003, 0.005, 0.007, 0.01], s: [0.1, 0.3, 0.5, 0.7, 0.9]
Complex Model Parameter Options	Fixed: g=0.15, q=0.000013
Round 1	k_0: [5, 20, 50, 100, 200, 300], p_{hi}: [0.05, 0.1, 0.2, 0.3, 0.5], s: [0.1, 0.3, 0.5, 0.7, 0.9],
Round 2	k_0: [150, 175, 200, 225, 250], p_{hi}: [0.10, 0.1125, 0.125, 0.1325, 0.15], s: [0.5, 0.6, 0.7, 0.8, 0.9]

Table 7: The list of values tested for each parameter, for each round for simulations of the full study period.

The other step I took to reduce the duration of the search process is that I tested for the best parameters using a single generated schedule of random logins; however, to assess model validity, I then ran each model with the best-performing set of parameters five times, with a new random login schedule each time, and reported the averaged results.

Model Results: All Geolocated Adopters

Next, I report the results of the simulations for the whole population of geolocated adopters. I find that for the entire study period, homophily and external factors can describe hashtag diffusion better than social contagion on Twitter. That being said, I find that online social contagion fueled the first forty-eight hours of the study period, though I am unable to distinguish between complex and simple contagion in my results.

First, I computed the best parameters for each model for the full population of geolocated adopters during the entire study period. Each cluster's cumulative proportional adoption rate was tightly correlated over the whole study period (as described in the previous chapter), therefore I did not conduct analysis on each cluster separately. The best-performing parameters and calculated least-squares cost are displayed in Table 8, with the least-squares averaged over five iterations of the model with a new random login schedule each time.

Model Type	Parameters	Mean Least-Squares Cost (of 5 Iterations)
Reference Model	q: 0.000013	472546
Simple Model	p: 0.0004 q: 0.000013 s: 0.9	391501
Complex Model	k_0: 300 p_{hi}: 0.15 g: 0.15 q: 0.000013 s: 0.8	635671

Table 8: The best parameters for each model, and the mean least-squares cost.

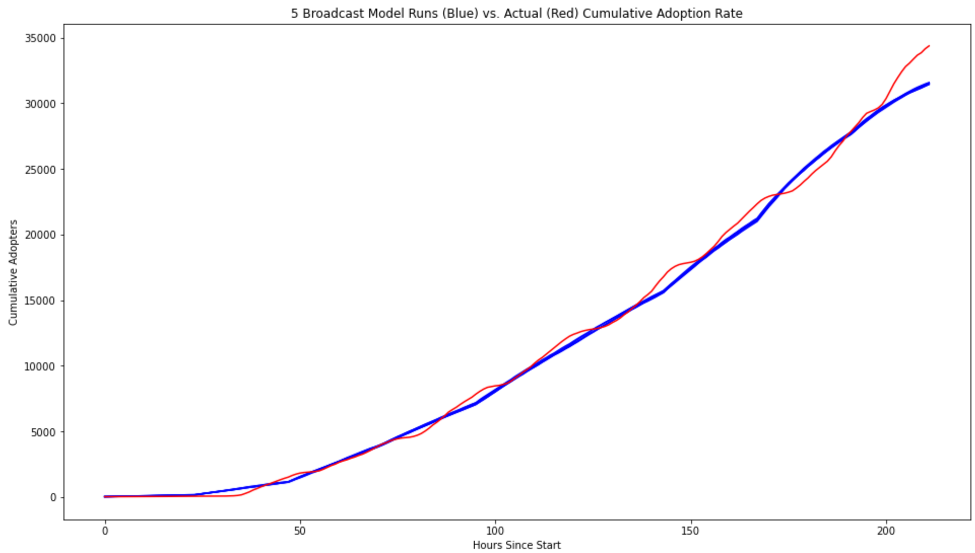

Figure 12: The reference model's predicted adoption curve performs extremely well compared to the actual adoption curve for the entire study period.

As is visible in Figure 12 and confirmed in the least-squares cost calculations, the reference model is a surprisingly strong illustration of the contagion rate, and effectively matches or beats the simple and complex models. This is despite the fact that the calculated *s*, the proportion of adoptions caused by external adoptions, for both the simple and complex models are above 70%; for the simple model, the best combination of parameters placed *s* at 0.9, and yet this model still only barely outperformed the reference model, which is likely due to the fact that a user's number of exposures do not increase as sharply overnight, so including exposure information may reduce the amount of overnight predicted adoptions (whereas the reference model does not make any assumptions about user diurnal tweet rates). The complex contagion model fairs slightly worse than the simple model, which is likely a result of having a higher number of parameters to tune. Both the complex and simple model performed better at higher values of *s*. It is striking that neither the complex nor simple model significantly improved upon the reference model, given that the reference model contains no information about individual agents' exposure within their network. This suggests that most of the perceived virality of #BlackLivesMatter in the week after George Floyd's death can actually be explained by users who happen to share ties being influenced by external factors.

Next, I repeated the steps of parameter search and model tuning, exclusively examining the first forty-eight hours of the study period. I re-used the same initial parameter options as defined earlier, but the second round of parameters were updated to reflect the range between the best two parameters when only considering the first forty-eight hours.

Upon completing my parameter search, I found that the models of contagion significantly outperformed the reference model for the first forty-eight hours of hashtag adoption. The mean least-squares cost for each model is included in Table 9 and the predicted adoption curves can be found in Figures 13, 14, and 15. Both the simple and complex models predict the exponential rise in adoptions that begins roughly thirty-four hours into the study period, while the reference model predicts a more linear rise in adoptions. This is evident in a much lower least-squares cost for both contagion models as compared to the reference model — the complex model's least-squares cost is less than one-tenth the cost of the reference model.

Model Type	Parameters	Mean Least-Squares Cost (of 5 Iterations)
Reference Model	q: 0.00001	49072
Simple Model	p: 0.0046 q: 0.00001 s: 0.1	7721
Complex Model	k_0: 50 p_{hi}: 0.08 g: 0.15 q: 0.00001 s: 0.15	4613

Table 9: The best performing parameters for each model for the first forty-eight hours.

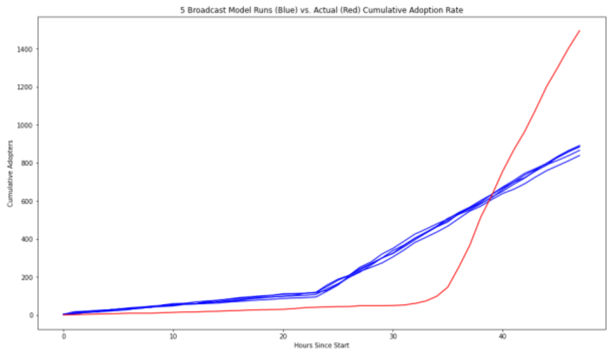

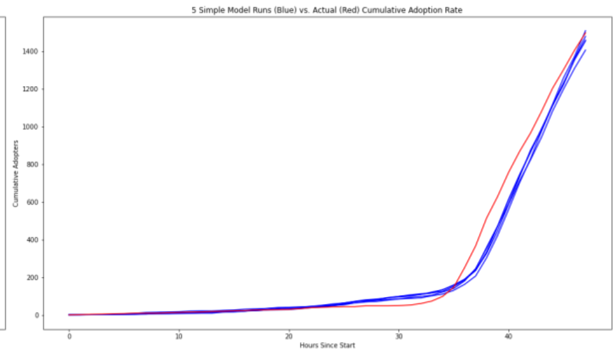

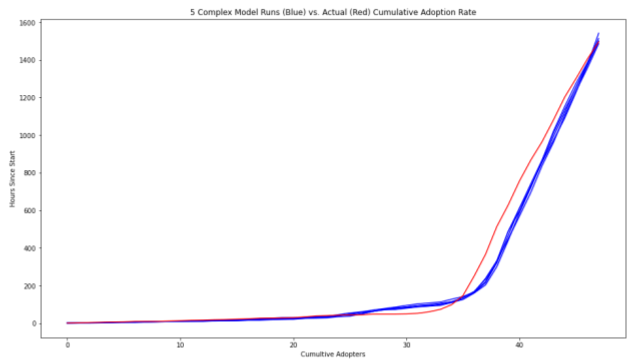

Figures 13, 14, and 15: The simulation results for each model for the first forty-eight hours.

An important caveat to this finding is that the reference model, by design, would be poorer at predicting shorter time periods, because it is based on daily news rate instead of hourly. However, I am confident in my results for two reasons: first, even when only considering the reference model's predicted adopters at the end of each full day, the reference model still overpredicts the number of adopters on the first day and underpredicts on the second day, demonstrating that the increase in news articles between days was insufficient to explain the rapid increase in Twitter users demonstrating their support for #BlackLivesMatter. Furthermore, over shorter periods of time, using the news rate to estimate discourse volume may be inaccurate given that users may not see an article immediately when it is posted, and news articles are not published evenly throughout the day.

Thus, during the earliest phase of adoption, social contagion on Twitter provided an impetus for users to demonstrate their support for #BlackLivesMatter, then, as the news cycle began to release more coverage of the movement, additional users logged in to Twitter to show public support. The earliest phase of adoption may be especially important in spurring journalists to write news articles, fueling a circular process by which media broadcasts amplify diffusion processes, a process which was suggested in (Freelon et al. 2018). Given that media coverage helps shape the political agenda (Edwards and Wood 1999; Hilgartner and Bosk 1988), this suggests that one major benefit of online activism is in directing the public conversation.

I do not draw a conclusion about which theoretical form of contagion best defines #BlackLivesMatter, because the simple and complex contagion models are both strong predictors of adoption behavior. Though the complex model performed slightly better than the simple model, as is evident in the complex model's cost of 4613 as compared to the simple model's cost of 7721, the shapes of their simulated adoption curves were very similar over the first forty-eight hours and it is unclear to what extent this distinction is analytically relevant. In order to limit my parameter search, I had fixed the lower-bound probability and the shape parameter for the complex contagion model. Had I trained on these parameters as well, I may have been able to achieve a better-fitting adoption curve for the complex model or been able to conclusively state that a model of simple contagion is as good at explaining hashtag diffusion. As computing power and the complexity of parameter search methods increase, future research should be able to achieve faster training and thus answer such questions more conclusively.

Regardless of which theoretical form of social contagion best describes #BlackLivesMatter, a high level of social reinforcement was necessary for early adoptions of #BlackLivesMatter. This is evident in the trained values of adoption probability after exposure.

As is illustrated in Figure 16, the simple model's best parameters indicate that a user had merely a 0.46% chance of adopting #BlackLivesMatter after a single exposure, while the complex model's best parameters indicate that a user had a maximum of an 8% chance of adoption after exposure, if at least fifty of the users they followed had already adopted. In either case, a user who adopted #BlackLivesMatter was likely exposed to many peers signaling support for the movement first. This finding would explain one of the key weaknesses of #BlackLivesMatter engagement — users are unlikely to engage with the hashtag outside of peaks of interest. If users only choose to engage in the conversations with the most volume, demonstrating long-term strength becomes challenging. I discuss the implications of this finding in greater depth in my discussion chapter, but before transitioning to my discussion, I must first conclude my contagion analysis by examining how model performance varies by cluster.

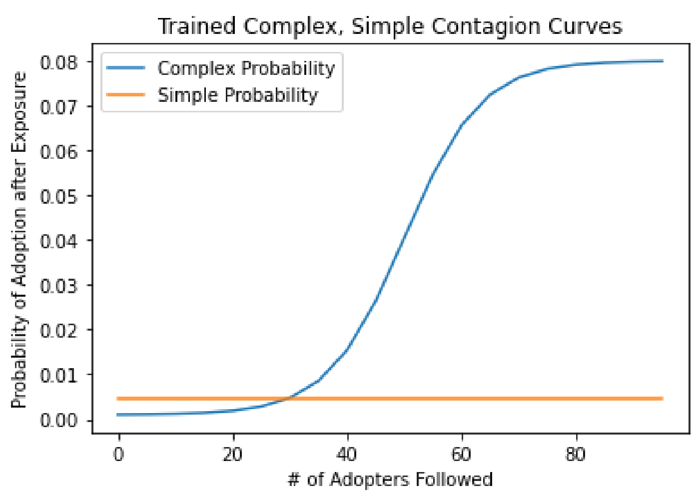

Figure 16: The probability of a user adopting #BlackLivesMatter after a single exposure, as calculated by the simple and complex contagion models.

Model Results: Analysis by Cluster

I had predicted that users from areas with higher rates of police violence and larger Black populations would require less social reinforcement to adopt #BlackLiveMatter than others. In order to test this hypothesis, I simulated user adoption for each cluster over the first forty-eight hours of the study period. First, I trained each model on each cluster separately, to determine if any model performed notably better for any cluster. Next, I computed the predicted adoption curves for each cluster using the model parameters as trained on the overall population, to facilitate easier comparison across clusters. I find that the simple and complex models each explain each cluster's behavior well, but that adopters from more diverse and urban county subdivisions required less social reinforcement to adopt #BlackLivesMatter than the overall population.

Upon training each model for each cluster, I found that a contagion model explains behavior well for all clusters. As shown in Table 10 of Appendix C, for every cluster, the reference model performed worse than the contagion models. Between the contagion models, however, performance was similar. The complex model performed slightly better for the MidCities, SubRur.Poorer, Sub.Richer, and Sub.Blacker clusters, while the simple model performed slightly better for the BigCities and Rur.White cluster, but these distinctions were only slight and not linked to any theoretical reasoning. For the purposes of my analysis, then, I claim that each contagion model performs equally well for predicting hashtag adoption. Having confirmed that there is not a drastic difference in contagion model performance across clusters, I transition to examining each cluster's simulated behavior using the parameters trained for the whole population of geolocated adopters.

When comparing model performance across clusters using a single set of parameters, it becomes evident that there are real differences in the level of social reinforcement required to adopt among adopters from different clusters. Given that my complex contagion model performed best overall, I include these results below in Table 11; the simple model results were similar. The Sub.Blacker cluster consistently had a higher number of adopters than their exposures on Twitter would predict. On the other hand, the only cluster which adopted at a notably later rate than the model would predict was the Rur.White cluster. This partially confirms my hypothesis that areas with high Black populations would require less social reinforcement for expressing support for #BlackLivesMatter, while whiter populations require more social reinforcement. The other cluster which adopted at a notably earlier time than the model would predict was the MidCities cluster, which, crucially, is the cluster containing Minneapolis and the areas around the city. In this case, I theorize that this result is due to unrecorded variations in local external influence. The daily news rate I used for constructing my reference model was based on the national news; however, given that the state of Minnesota saw widespread Black Lives Matter protests before the rest of the county (Taylor 2021), users from this cluster likely had a higher level of external influence within the first forty-eight hours of the study period than the national news rate would suggest. It may also be the case that users from MidCities were especially motivated to engage with #BlackLivesMatter given the local significance of Floyd’s murder. Distinguishing between these motivations would likely require a qualitative examination of user’s tweets, so I leave this question to further research, but note that I had identified early adoptions by Sub.Blacker and MidCities adopters as having special importance in the diffusion process in the previous chapter, so this finding helps explain why individuals from these areas adopted at an earlier rate — they did not wait for as much social reinforcement online to signal their support for Black Lives Matter.

Surprisingly, the BigCities cluster did not similarly exhibit a lower level of requisite social reinforcement, even though the BigCities cluster was the only cluster to adopt at a higher per capita rate than the Sub.Blacker cluster. This suggests that BigCities users were being exposed to #BlackLivesMatter at higher rates than other users, because otherwise, a high probability of adoption after each exposure would be necessary for these individuals to adopt #BlackLivesMatter at a much higher per capita rate than other cluster. This leads me to theorize that part of the surge in adoption from the BigCities cluster comes from that population's higher Twitter usage overall in comparison to the other clusters. This finding does contrast with an argument I made in the previous chapter, where I noted that higher Twitter usage among urban residents was not enough to explain their disproportionately high per capita adoption rate. It may be the case, then, that Twitter's reputation as a site for activism (Tillery 2019) attracted urban supporters of the movement to the platform at higher rates. The BigCities cluster exhibited a high rate of previous adoption (34.6%) — the percentage of users who had used #BlackLivesMatter before the study period. Only the Sub.Blacker cluster exhibits a higher rate of previous adoption (35.8%) and all other clusters exhibit a previous adoption rate of less than 30%. Understanding who was adopting Black Lives Matter within cities is an important question which I leave to further research, as I note in my discussion.

Predicted (Complex Model) vs. Empirical Adoption Curve	
Rur.White 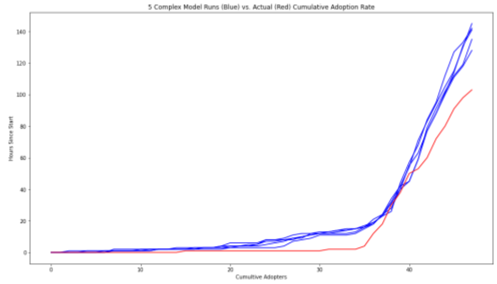	BigCities 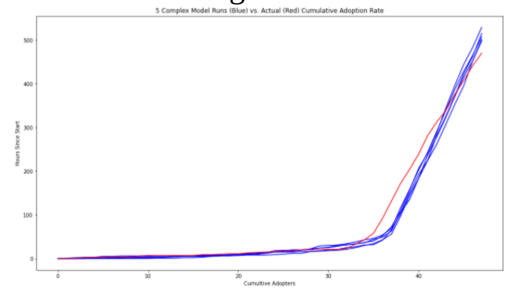
Sub.Blacker 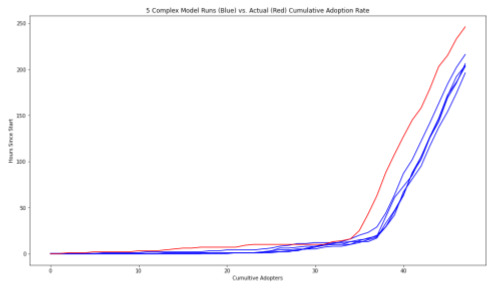	SubRur.Poorer 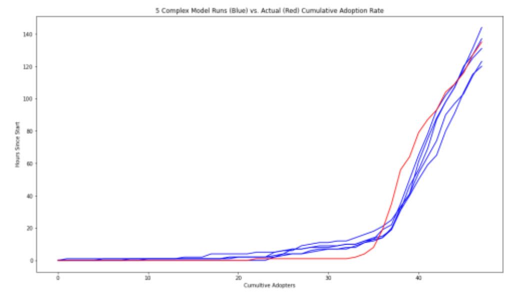
Sub.Richer 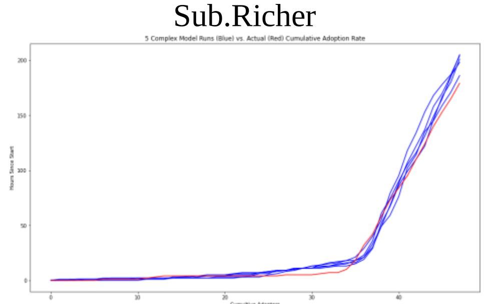	MidCities 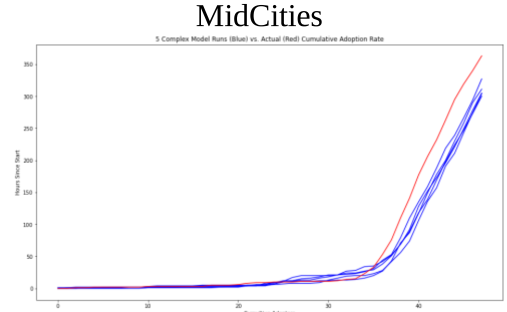

<u>Table 11: The predicted adoptions for each cluster using the complex model and the parameters trained on the total geolocated population.</u>

Overall, these findings reinforce themes previously identified in the prior chapter. We previously identified that BigCities adopters were central in the #BlackLivesMatter geolocated network; this fact now helps explain why they adopted at such a high rate. MidCities and Sub.Blacker adopters' importance in the earliest stage of hashtag diffusion is reinforced, as these users adopted at lower levels of social reinforcement, which may have in-turn influenced other users to adopt sooner. Finally, Rur.White subdivisions' low rate of per capita adoption is at least partially explained by Rur.White users' need for more social reinforcement before adopting #BlackLivesMatter.

Summarizing Findings

Now I answer my final research question, and assess my final two hypotheses:

RQ-3: How can social contagion explain the diffusion of #BlackLivesMatter?

H3.1: The #BlackLivesMatter hashtag is a complex contagion, requiring multiple exposures before adoption, but the contagion effect will be moderated when accounting for homophily among users.

H3.2: Users in more urban areas with higher rates of police violence and larger Black populations will require less social reinforcement to adopt #BlackLiveMatter han other users.

I found that the diffusion of #BlackLivesMatter over the entire study period can largely be explained by homophily and external influences such as traditional news sources. The first forty-eight hours of adoption, however, are best explained by a process of online social contagion. Both simple and complex models perform well at predicting the rate of adoption during this period. Regardless of which type of contagion #BlackLivesMatter represents, the high level of social reinforcement necessary for adoption can help explain one of the ongoing features of #BlackLivesMatter — its usage is episodic, characterized by high peaks and long valleys (Freelon et al. 2016). To understand why this would be the case, I return to a concept from my literature review — liveness, a feature of Twitter that makes users feel they “need to contribute to conversations as they happen” (Zulli 2020). Liveness means users focus on trending events but may neglect topics once they drop out of the spotlight. For a social movement interested in signaling movement capacity, these volatile temporal dynamics can be a hindrance to movement success, as politicians learn that they can wait for a topic to stop trending rather than enact policy change.

When examining hashtag adoption by cluster, I found that whiter, more rural areas required higher levels of social reinforcement to adopt #BlackLivesMatter than other areas, confirming my second hypothesis. Beyond reinforcing the themes described in my previous chapter, this result is relevant for two reasons: first, it connects individual decisions about racial issues to the macro-social process in which a movement attempts to alter the racial power structure. For sociologists interested in connecting micro- and macro-levels of analysis, contagion theories can provide that link. This result is also relevant because it highlights how contagion studies can benefit by taking into account an individual's social identities. Both contagion and social identity have been used to explain how social movements grow, but rarely have they both been addressed in the same analysis (for an exception to this statement, see: State and Adamic 2015). When we do combine them, we can learn more about why and when racial hierarchies are questioned.

CHAPTER 7: DISCUSSION

Results Summary

In this thesis, I tackled several research questions:

RQ1: How did the #BlackLivesMatter hashtag diffuse on Twitter in the aftermath of George Floyd's death?

First, I found that outside of the peak periods of hashtag usage, #BlackLivesMatter is maintained by a dense network of highly-active Twitter users. As news of George Floyd's death spread, less-active and less-connected users also began to adopt #BLACKLIVESMATTER. Users outside the United States played a significant role in driving hashtag usage, making this phenomena a truly global diffusion event and contributing to the unprecedented scale of movement activity.

RQ2: How did hashtag adoption vary by neighborhood characteristics?

Next, I examined hashtag diffusion stratified by county subdivision characteristics. I found that metrics of racial injustice were less informative than the overall demographic patterns that have characterized the country's recent political divisions. Population density was the most important predictor of county subdivision adoption rates, and individuals from urban areas were the most likely to adopt #BlackLivesMatter on Twitter, as well as the most densely connected users. Thus, while this was a global and nationwide event, individuals from cities drove the conversation. My hypothesis for this question was only partially correct, as I expected police violence and racial income inequality to be more salient for hashtag diffusion rates, while my expectation that population density and racial demographics would be salient proved true.

RQ3: How can social contagion explain the diffusion of #BlackLivesMatter?

Finally, I found that while online social contagion is a poor predictor of hashtag diffusion over the whole study period, online social contagion plays an important role in the earliest adoptions, which may then drive a broader cultural conversation. I had hypothesized that #BlackLivesMatter would represent a complex contagion, and that this effect would be moderated when accounting for homophily and external factors. Instead, I found that when accounting for external factors, homophily is a very strong explanation for perceived contagion on Twitter. Still, contagion models work well when only examining the first 48 hours of my study period, before widespread news coverage of George Floyd's murder. While I do not examine a causal relationship between contagion and news in this study, one mechanism to explain my findings would be that early contagion fuels news coverage, which in turns fuels more hashtag adoption. I also found that county subdivision characteristics influenced the level of social reinforcement necessary for adoption, with individuals in more rural and whiter areas requiring more reinforcement to adopt than other users.

Implications

To understand the implications of my results, I return to several concepts which I outlined in my literature review. In my literature review, I defined Black Lives Matter as a hybrid movement, engaging both the logics of collective action — traditional top-down organizing — and connective action — action spreading through "large-scale, fluid social networks" without strong collective identity (Bennett and Sederberg 2012:748). Within the broader movement, online #BlackLivesMatter engagement represented a form of connective action, where the hashtag itself represented a counterpublic that users could participate in and question dominant narratives. Now I expand upon these concepts with several of the primary implications of my research.

First, the dominance of the BigCities cluster points to the heterogeneity of the #BlackLivesMatter counterpublic. This finding could be interpreted in two ways: on the one hand, the dominance of the BigCities cluster may be better attributed to the modern urban-rural political divide in this country, along with the age-old dynamic whereby those individuals from cities dominate the discourse. Alternatively, these urban voices could represent marginalized individuals living in the most concentrated regions of police violence — those individuals who would previously not have had the ability to influence the public sphere but who are now able to. This utopian ideal is suggested in (Welles and Jackson 2019), where they find that Twitter provides an opportunity for urban individuals to counter mainstream narratives in the media about urban unrest. It is notable that urban areas fueled hashtag diffusion, given that the construction of the modern American city was fueled by racialized policies on everything from housing to infrastructure and education, White flight, and broken-windows policing (Derickson 2017). After all, without racial inequality, there would not be a predominantly white Sub.Richer cluster or a diverse BigCities cluster. While further research should be able to conclusively argue what urban engagement represents, what my research does reveal is that urban individuals are essential to the process of hashtag diffusion.

Second, my research demonstrates some of the strengths and weaknesses of connective action. Many have wondered now that we are two years after Floyd's murder whether the widespread activism of June 2020 had any impact. Especially in the aftermath of the racially-motivated shooting in Buffalo and at a time when Black Americans are feeling more pessimistic than ever about racism (Foster-Frau et al. 2022), it is natural to ask if digitally-fueled activism matters. Given that my research focused on understanding the mechanisms of hashtag diffusion rather than its impact, I mostly leave this question to future research, but note a few ways that my research can inform this debate.

Consider, for example, the speed with which Black Lives Matter protests spread. Nationwide protests began within two days of George Floyd’s murder. By contrast, the March on Washington took ten years to plan (Tufekci 2020). One could argue whether or not protests would have spread as quickly without Twitter, but Twitter has often facilitated other offline protests, and news of Floyd’s death was circulating widely on Twitter within twenty-four hours of the event, so this seems likely (Jost et al. 2018).

In addition, another important function of #BlackLivesMatter, as revealed in my network analysis, is that the hashtag linked users who otherwise inhabited very different geographies. Even individuals living in predominantly White, rural areas were receiving information directly from individuals in cities where major protests were occuring and from across the country. While physical protests do bring together a variety of people, social media allows for individuals to have an even more diverse sphere of influence, at least geographically.

Finally, the high level of social reinforcement I identified as necessary for engagement does help explain what others have noted as a flaw of online activism: that it is characterized by intense periods of engagement but not alway sustaining of a movement. #BlackLivesMatter is hardly the only form of online activism to be described in this way; #MeToo has also been noted as struggling with this phenomena (Lindgren 2019). In the previous chapter, I noted how the concept of liveness can help explain why this occurs. For a movement seeking to demonstrate capacity, this feature of online activism is a major hurdle. Future activists should direct their attention towards identifying how to sustain engagement outside of peak periods, and furthermore, how to demonstrate to the political elites that this engagement is meaningful.

Beyond informing the debate over the efficacy and function of hashtag activism, an additional implication of my research regards the methodological innovations I have introduced to social contagion research. Much of the social contagion and online social networks literature conceives of users as identity-less blobs, existing only online without their offline social realities influencing their behavior, or across a single salient dimension, like political orientation (Wang et al. 2020). Geolocation is an imperfect technique for stratifying users by salient characteristics, yet this basic stratification has already revealed significant differences in online behavior driven by offline variation. Future researchers can relatively easily apply this method to consider user identity in their research. Furthermore, my alterations of the contagion models of Fink et al. (2016) and my stratification of user populations by geography in my contagion simulations address the issue of homophily in a theoretically-driven manner, and as far as I know, no other contagion simulation research has attempted this. All of these methodologies describe ways in which the sheer scale of online networks data can be combined with other relatively-easily available metrics in order to paint a more accurate picture of diffusion dynamics. As the agent-based modeling field continues to mature, I expect future researchers to continue developing models which better reflect social realities without the loss of analytical usefulness.

Limitations and Future Research

My research has several limitations, each of which pose a potential route for future research.

First, as with any observational study using social media data, my data collection was shaped by what data social media platforms choose to make available to researchers. I identified some ways to adapt to the limitations of Twitter's API, for example by examining geographic location as a method of considering users' salient characteristics, but I was not able to address every shortcoming. For example, the Twitter API only provides access to a user's current followers list, not their historical followers, thus the network I examined is likely different than the true #BlackLivesMatter followers network. Future research should examine how follower-networks evolve, which could allow researchers to estimate how different their network is from the true network and how urgent data collection in the aftermath of a major diffusion event is if researchers want to conduct accurate research.

Additionally, as mentioned in my literature review, platform data limitations not only shape research methodology, but they also shape which platforms get studied at all; Twitter's relatively high level of access has made it a favorite site of inquiry for researchers. Slowly, researchers have been identifying ways of studying other platforms. For example, Chang, Richardson, and Ferrara (2021) detailed the spread of Black Lives Matter on Instagram, finding that the platform was a critical source of protest journalism in the two weeks after George Floyd's death. Conducting research on other platforms may be costly and time-consuming but is essential for understanding online activism.

Another limitation was that, while my approach of geotagging was informative, I was unable to examine who within a county subdivision was engaging with the movement. This shortcoming points for the need for mixed methods in social media research. Studies which poll individual users about their engagement could answer this question. Extending this line of thought further, mixed methods could also deepen our understanding of the contagion processes I outline in this research. Qualitative studies, for example, asking users why they made the decision to engage with a movement when they did, could help explain how activists can overcome the “bursty” nature of online activism to form more durable movements.

Finally, I note the need to further validate my modeling approach. Any model of human behavior requires assumptions and simplifications; the goal of the researcher is to identify the model which most closely predicts behavior while also being simple enough to be analytically useful. My contagion model design in Chapter 6 was shaped by limitations of prior contagion studies, but further testing these models on additional hashtags would provide further evidence as to their validity while also allowing future contagion research to account for the existence of homophily, which was a major confounding factor in my analysis.

Overall, I hope my research can provide context for the diffusion event in the aftermath of George Floyd’s death, which can serve as a foundation for future research to explore this event and this movement more deeply. It is important not only to understand how online activism works, but also to understand how it can enact social change, and that is a question future researchers must continue to ask.

WORKS CITED

Alfano, Mark, Ritsaart Reimann, Ignacio Quintana, Marc Cheong, and Colin Klein. 2021. "The Affiliative Use of Emoji and Hashtags in the Black Lives Matter Movement: A Twitter Case Study."

Alfonzo, Paige. 2021. "A Topology of Twitter Tactics: Tracing the Rhetorical Dimensions and Digital Labor of Networked Publics." *Social Media + Society* 7(2):20563051211025510. doi: 10.1177/20563051211025514.

Auxier, Brooke, and Monica Anderson. 2021. "Social Media Use in 2021." *Pew Research Center: Internet, Science & Tech*. Retrieved May 17, 2022 (https://www.pewresearch.org/internet/2021/04/07/social-media-use-in-2021/).

Bakshy, Eytan, Itamar Rosenn, Cameron Marlow, and Lada Adamic. 2012. "The Role of Social Networks in Information Diffusion." Pp. 519–28 in *Proceedings of the 21st international conference on World Wide Web, WWW '12*. New York, NY, USA: Association for Computing Machinery.

Barash, Vladimir. 2011. "The Dynamics Of Social Contagion." PhD dissertation, Information Science, Cornell University.

Barash, Vladimir, and John Kelly. 2012. *Salience vs. Commitment: Dynamics of Political Hashtags in Russian Twitter*. 2012–9. The Berkman Center for Internet & Society.

Bennett, Lance, and Alexandra Segerberg. 2012. "The Logic of Connective Action." *Information, Communication & Society* 15:1–30. doi: 10/f2znbx.

Benton, Joshua. 2019. “Twitter Is Removing Precise-Location Tagging on Tweets — a Small Win for Privacy but a Small Loss for Journalists and Researchers.” *Nieman Lab*. Retrieved October 19, 2021 (https://www.niemanlab.org/2019/06/twitter-is-turning-off-location-data-on-tweets-a-small-win-for-privacy-but-a-small-loss-for-journalists-and-researchers/).

Berry, George, Christopher J. Cameron, Patrick Park, and Michael W. Macy. 2017. “The Opacity Problem in Social Contagion.” *Social Networks* 56:93–101. doi: 10/ghjn6m.

Borge-Holthoefer, Javier, Sandra González-Bailón, Alejandro Rivero, and Yamir Moreno. 2014. “The Spanish ‘Indignados’ Movement: Time Dynamics, Geographical Distribution, and Recruitment Mechanisms.” Pp. 155–77 in *Online Collective Action: Dynamics of the Crowd in Social Media*, *Lecture Notes in Social Networks*, edited by N. Agarwal, M. Lim, and R. T. Wigand. Vienna: Springer.

Brock, André. 2012. “From the Blackhand Side: Twitter as a Cultural Conversation.” *Journal of Broadcasting & Electronic Media* 56(4):529–49. doi: 10/gddxwm.

Brown, Jennings. 2015. “A Year After Michael Brown, Black Lives Still Matter.” *Vocativ*. Retrieved September 24, 2021 (https://www.vocativ.com/218365/michael-brown-and-black-lives-matter/).

Bruns, Axel, and Tim Highfield. 2016. “Is Habermas on Twitter? Social Media and the Public Sphere.” Pp. 56–73 in *The Routledge companion to social media and politics*, edited by G. Enli, A. Bruns, A. O. Larsson, E. Skogerbo, and C. Christensen. United States of America: Routledge.

Bruns, Axel, and Hallvard Moe. 2014. “Structural Layers of Communication on Twitter.” Pp. 15–28 in *Twitter and society*. Vol. 89, *Digital Formations*, edited by A. Bruns, M. Mahrt, K. Weller, J. Burgess, and C. Puschmann. United States of America: Peter Lang Publishing.

Buchanan, Larry, Quoctrung Bui, and Jugal K. Patel. 2020. “Black Lives Matter May Be the Largest Movement in U.S. History.” *The New York Times*, July 3.

Bureau, US Census. 2020. "American Community Survey (ACS)." *Census.Gov*. Retrieved June 6, 2022 (https://www.census.gov/programs-surveys/acs).

Campaign Zero. 2022. "Mapping Police Violence." *Mapping Police Violence*. Retrieved June 6, 2022 (https://mappingpoliceviolence.org/).

Carney, Nikita. 2016. "All Lives Matter, but so Does Race: Black Lives Matter and the Evolving Role of Social Media." *Humanity & Society* 40(2):180–99. doi: 10/gf3gsp.

Chang, Ho-Chun Herbert, Allissa Richardson, and Emilio Ferrara. 2021. "#JusticeforGeorgeFloyd: How Instagram Facilitated the 2020 Black Lives Matter Protests."

Chase, Garrett. 2017. "The Early History of the Black Lives Matter Movement, and the Implications Thereof Notes." *Nevada Law Journal* 18(3):1091–1112.

Cobb, Jelani. 2016. "The Matter of Black Lives." *The New Yorker*, March 6.

Conover, Michael, Jacob Ratkiewicz, Matthew Francisco, Bruno Goncalves, Filippo Menczer, and Alessandro Flammini. 2011. "Political Polarization on Twitter." *Proceedings of the International AAAI Conference on Web and Social Media* 5(1):89–96.

De Choudhury, Munmun, Shagun Jhaver, Benjamin Sugar, and Ingmar Weber. 2016. "Social Media Participation in an Activist Movement for Racial Equality." *Proceedings of the ... International AAAI Conference on Weblogs and Social Media. International AAAI Conference on Weblogs and Social Media* 2016:92–101.

De Longueville, Bertrand, Robin S. Smith, and Gianluca Luraschi. 2009. "'OMG, from Here, I Can See the Flames!': A Use Case of Mining Location Based Social Networks to Acquire Spatio-Temporal Data on Forest Fires." Pp. 73–80 in *Proceedings of the 2009 International Workshop on Location Based Social Networks*, *LBSN '09*. New York, NY, USA: Association for Computing Machinery.

Dennis, Rutledge, and Kimya Dennis. 2020. “Confrontational Politics: The Black Lives Matter Movement.” Pp. 13–27 in *The Wiley Blackwell Companion to Race, Ethnicity, and Nationalism*. John Wiley & Sons, Ltd.

Derickson, Kate Driscoll. 2017. “Urban Geography II: Urban Geography in the Age of Ferguson.” *Progress in Human Geography* 41(2):230–44. doi: 10.1177/0309132515624315.

Diani, Mario. 1992. “The Concept of Social Movement.” *The Sociological Review* 40(1):1–25. doi: 10.1111/j.1467-954X.1992.tb02943.x.

Dixon, Marc, and Vincent J. Roscigno. 2003. “Status, Networks, and Social Movement Participation: The Case of Striking Workers.” *American Journal of Sociology* 108(6):1292–1327. doi: 10.1086/375198.

Edrington, Candice LaShara, and Nicole Lee. 2018. “Tweeting a Social Movement: Black Lives Matter and Its Use of Twitter to Share Information, Build Community, and Promote Action.” *The Journal of Public Interest Communications* 2(2):289–306. doi: 10.32473/jpic.v2.i2.p289.

Edwards, George C., and B. Dan Wood. 1999. “Who Influences Whom? The President, Congress, and the Media.” *The American Political Science Review* 93(2):327–44. doi: 10.2307/2585399.

Fink, Clay, Aurora Schmidt, Vladimir Barash, Christopher Cameron, and Michael Macy. 2015. “Complex Contagions and the Diffusion of Popular Twitter Hashtags in Nigeria.” *Social Network Analysis and Mining* 6(1). doi: 10.1007/s13278-015-0311-z.

Fink, Clay, Aurora Schmidt, Vladimir Barash, John Kelly, Christopher Cameron, and Michael Macy. 2016. “Investigating the Observability of Complex Contagion in Empirical Social Networks.” Pp. 121–30 in *Proceedings of the Tenth International AAAI Conference on Web and Social Media*. Association for the Advancement of Artificial Intelligence.

Foster-Frau, Silvia, Arelis Hernández, Scott Clement, and Emily Guskin. 2022. "Poll: Black Americans Fear More Racist Attacks after Buffalo Shooting." *Washington Post*, May 21.

Fraser, Nancy. 1990. "Rethinking the Public Sphere: A Contribution to the Critique of Actually Existing Democracy." *Social Text* (25/26):56–80. doi: 10/fgp8qj.

Freelon, Deen, Charlton McIllwain, and Meredith Clark. 2016. *Beyond the Hashtags: #Ferguson, #Blacklivesmatter, and the Online Struggle for Offline Justice*. Center for Media and Social Impact.

Freelon, Deen, Charlton McIlwain, and Meredith Clark. 2018. "Quantifying the Power and Consequences of Social Media Protest." *New Media & Society* 20(3):990–1011. doi: 10.1177/1461444816676646.

Garza, Alicia, and L. A. Kauffman. 2015. "A Love Note to Our Folks." *N+1*. Retrieved October 25, 2021 (https://www.nplusonemag.com/online-only/online-only/a-love-note-to-our-folks/).

Gladwell, Malcolm. 2010. "Small Change." *The New Yorker*, September 27.

Granovetter, Mark S. 1973. "The Strength of Weak Ties." *American Journal of Sociology* 78(6):1360–80. doi: 10.1086/225469.

Greijdanus, Hedy, Carlos A. de Matos Fernandes, Felicity Turner-Zwinkels, Ali Honari, Carla A. Roos, Hannes Rosenbusch, and Tom Postmes. 2020. "The Psychology of Online Activism and Social Movements: Relations between Online and Offline Collective Action." *Current Opinion in Psychology* 35:49–54. doi: 10.1016/j.copsyc.2020.03.003.

Guilbeault, Douglas, Joshua Becker, and Damon Centola. 2018. "Complex Contagions: A Decade in Review." Pp. 3–25 in *Complex Spreading Phenomena in Social Systems: Influence and Contagion in Real-World Social Networks*, edited by S. Lehmann and Y.-Y. Ahn. Cham: Springer International Publishing.

Habermas, Jürgen. 1989. *The Structural Transformation of the Public Sphere: An Inquiry into a Category of Bourgeois Society*. Cambridge, Mass. : MIT Press.

Halberstam, Yosh, and Brian Knight. 2016. "Homophily, Group Size, and the Diffusion of Political Information in Social Networks: Evidence from Twitter." *Journal of Public Economics* 143:73–88. doi: 10/f88stf.

Hanappi, Hardy. 2017. "Agent-Based Modelling History, Essence, Future." *PSL Quarterly Review* 70.

Hasan, Samiul, Xianyuan Zhan, and Satish V. Ukkusuri. 2013. "Understanding Urban Human Activity and Mobility Patterns Using Large-Scale Location-Based Data from Online Social Media." P. 1 in *Proceedings of the 2nd ACM SIGKDD International Workshop on Urban Computing - UrbComp '13*. Chicago, Illinois: ACM Press.

Hatmaker, Taylor. 2021. "Facebook's New Academic Research API Opens in Early Access." *TechCrunch*. Retrieved June 2, 2022 (https://techcrunch.com/2021/11/15/facebooks-researcher-api-meta-academic-research/).

Hecht, Brent, Lichan Hong, Bongwon Suh, and Ed H. Chi. 2011. "Tweets from Justin Bieber's Heart: The Dynamics of the 'Location' Field in User Profiles." Pp. 237–46 in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*.

Hilgartner, Stephen, and Charles L. Bosk. 1988. "The Rise and Fall of Social Problems: A Public Arenas Model." *American Journal of Sociology* 94(1):53–78. doi: 10.1086/228951.

Hill, Marc Lamont. 2018. "'Thank You, Black Twitter': State Violence, Digital Counterpublics, and Pedagogies of Resistance." *Urban Education* 53(2):286–302. doi: 10/gjqtnw.

Hobson, Elizabeth A., Matthew J. Silk, Nina H. Fefferman, Daniel B. Larremore, Puck Rombach, Saray Shai, and Noa Pinter-Wollman. 2021. "A Guide to Choosing and Implementing Reference Models for Social Network Analysis." *Biological Reviews* 96(6):2716–34. doi: 10/gk36kx.

Hodas, Nathan O., and Kristina Lerman. 2014. "The Simple Rules of Social Contagion." *Scientific Reports* 4(1). doi: 10/ggvr2q.

Hu, Yingjie, and Ruo-Qian Wang. 2020. "Understanding the Removal of Precise Geotagging in Tweets." *Nature Human Behaviour* 4(12):1219–21. doi: 10.1038/s41562-020-00949-x.

Huang, Xiao, Zhenlong Li, Yuqin Jiang, Xiaoming Li, and Dwayne Porter. 2020. "Twitter Reveals Human Mobility Dynamics during the COVID-19 Pandemic" edited by S. Gao. *PLOS ONE* 15(11):e0241957. doi: 10.1371/journal.pone.0241957.

Iceland, John, Daniel H. Weinberg, and Erika Steinmetz. 2002. *Racial and Ethnic Residential Segregation in the United States 1980-2000*. Vol. 8. U.S. Census Bureau.

Ince, Jelani, Fabio Rojas, and Clayton A. Davis. 2017. "The Social Media Response to Black Lives Matter: How Twitter Users Interact with Black Lives Matter through Hashtag Use." *Ethnic and Racial Studies* 40(11):1814–30. doi: 10.1080/01419870.2017.1334931.

Jackson, Sarah J., and Brooke Foucault Welles. 2016. "#Ferguson Is Everywhere: Initiators in Emerging Counterpublic Networks." *Information,Communication&Society* 19(3):397–418. doi: 10.1080/1369118X.2015.1106571.

Jiang, Julie, Emily Chen, Shen Yan, Kristina Lerman, and Emilio Ferrara. 2020. "Political Polarization Drives Online Conversations about COVID-19 in the United States." *Human Behavior and Emerging Technologies* 2(3):200–211. doi: 10.1002/hbe2.202.

Jost, John T., Pablo Barberá, Richard Bonneau, Melanie Langer, Megan Metzger, Jonathan Nagler, Joanna Sterling, and Joshua A. Tucker. 2018. "How Social Media Facilitates Political Protest: Information, Motivation, and Social Networks." *Political Psychology* 39(S1):85–118. doi: 10/gc3m4d.

Khanam, Kazi Zainab, Gautam Srivastava, and Vijay Mago. 2022. “The Homophily Principle in Social Network Analysis: A Survey.” *Multimedia Tools and Applications*. doi: 10.1007/s11042-021-11857-1.

Latour, Bruno. 2005. *Reassembling the Social: An Introduction to Actor-Network-Theory*. Oxford ; New York: Oxford University Press.

Law, John. 1992. “Notes on the Theory of the Actor-Network: Ordering, Strategy, and Heterogeneity.” *Systems Practice* 5(4):379–93. doi: 10/c735j4.

Lerman, Kristina. 2016. “Information Is Not a Virus, and Other Consequences of Human Cognitive Limits.” *Future Internet* 8(2):21. doi: 10.3390/fi8020021.

Lindgren, Simon. 2019. “Movement Mobilization in the Age of Hashtag Activism: Examining the Challenge of Noise, Hate, and Disengagement in the #MeToo Campaign.” *Policy & Internet* 11(4):418–38. doi: 10.1002/poi3.212.

Lloyd, S. 1982. “Least Squares Quantization in PCM.” *IEEE Transactions on Information Theory* 28(2):129–37. doi: 10.1109/TIT.1982.1056489.

Lupia, Arthur, and Gisela Sin. 2003. “Which Public Goods Are Endangered?: How Evolving Communication Technologies Affect The Logic of Collective Action.” *Public Choice* 117(3–4):315–31. doi: 10.1023/b:puch.0000003735.07840.c7.

Makalani, Minkah. 2017. “Black Lives Matter and the Limits of Formal Black Politics.” *South Atlantic Quarterly* 116(3):529–52. doi: 10.1215/00382876-3961472.

Martineau, Paris. 2019. “You May Have Forgotten Foursquare, but It Didn’t Forget You.” *Wired*, March 8.

Marwick, Alice E., and danah boyd. 2011. “I Tweet Honestly, I Tweet Passionately: Twitter Users, Context Collapse, and the Imagined Audience.” *New Media & Society* 13(1):114–33. doi: 10.1177/1461444810365313.

Massey, Douglas S., Jonathan Rothwell, and Thurston Domina. 2009. “The Changing Bases of Segregation in the United States.” *The ANNALS of the American Academy of Political and Social Science* 626(1):74–90. doi: 10/cnrspf.

McAdam, Doug. 1986. “Recruitment to High-Risk Activism: The Case of Freedom Summer.” *American Journal of Sociology* 92(1):64–90. doi: 10/fngpg9.

McDermott, Monica, and FrankL. Samson. 2005. “White Racial and Ethnic Identity in the United States.” *Annual Review of Sociology* 31(1):245–61. doi: 10.1146/annurev.soc.31.041304.122322.

McPherson, Miller, Lynn Smith-Lovin, and James M. Cook. 2001. “Birds of a Feather: Homophily in Social Networks.” *Annual Review of Sociology* 27(1):415–44. doi: 10/fs9tff.

Melucci, Alberto. 1995. “The Process of Collective Identity.” Pp. 41–63 in *Social Movements and Culture*. University of Minnesota Press.

Meta, Inc. n.d. “Instagram Platform Documentation.” *Meta for Developers*. Retrieved June 2, 2022 (https://developers.facebook.com/docs/instagram/).

Mønsted, Bjarke, Piotr Sapieżyński, Emilio Ferrara, and Sune Lehmann. 2017. “Evidence of Complex Contagion of Information in Social Media: An Experiment Using Twitter Bots.” *PLOS ONE* 12(9):e0184148. doi: 10/gbx6vc.

Moore, Darnell L., and Patrisse Cullors. 2014. “5 Ways to Never Forget Ferguson – and Deliver Real Justice for Michael Brown.” *The Guardian*, September 4.

Mundt, Marcia, Karen Ross, and Charla M. Burnett. 2018. "Scaling Social Movements Through Social Media: The Case of Black Lives Matter." *Social Media + Society* 4(4):2056305118807911. doi: 10.1177/2056305118807911.

Nexis Uni. n.d. "Nexis Uni® Home." *Nexis Uni*. Retrieved June 5, 2022 (https://advance-lexis-com.dartmouth.idm.oclc.org/bisacademicresearchhome?crid=84e21e92-3f10-4b5f-b6c1-ac76eeb90075&pdmfid=1516831&pdisurlapi=true).

Ogburn, Elizabeth L. 2018. "Challenges to Estimating Contagion Effects from Observational Data." Pp. 47–64 in *Complex Spreading Phenomena in Social Systems: Influence and Contagion in Real-World Social Networks*, *Computational Social Sciences*, edited by S. Lehmann and Y.-Y. Ahn. Cham: Springer International Publishing.

Olson, Mancur. 1965. *The Logic of Collective Action ; Public Goods and the Theory of Groups.* Cambridge, Mass: Harvard University Press.

Olteanu, Alexandra, Ingmar Weber, and Daniel Gatica-Perez. 2015. "Characterizing the Demographics Behind the #BlackLivesMatter Movement." in *Proceedings of the AAAI Spring Symposia on Observational Studies through Social Media and Other Human-Generated Content*. Stanford, US.

Parks, Virginia. 2012. "The Uneven Geography of Racial and Ethnic Wage Inequality: Specifying Local Labor Market Effects." *Annals of the Association of American Geographers* 102(3):700–725. doi: 10.1080/00045608.2011.600198.

Pedregosa, F., G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. 2011. "Scikit-Learn: Machine Learning in Python." *Journal of Machine Learning Research* 12:2825–30.

Rickford, Russell. 2016. “Black Lives Matter: Toward a Modern Practice of Mass Struggle.” *New Labor Forum* 25(1):34–42. doi: 10.1177/1095796015620171.

Romero, Daniel M., Brendan Meeder, and Jon Kleinberg. 2011. “Differences in the Mechanics of Information Diffusion across Topics: Idioms, Political Hashtags, and Complex Contagion on Twitter.” Pp. 695–704 in *Proceedings of the 20th international conference on World wide web, WWW ’11*. New York, NY, USA: Association for Computing Machinery.

scikit-learn developers. n.d. “Clustering.” *Scikit Learn*. Retrieved April 26, 2022 (https://scikit-learn/stable/modules/clustering.html).

Shalizi, Cosma Rohilla, and Andrew C. Thomas. 2011. “Homophily and Contagion Are Generically Confounded in Observational Social Network Studies.” *Sociological Methods & Research* 40(2):211–39. doi: 10.1177/0049124111404820.

Sharma, Sanjay. 2013. “Black Twitter? Racial Hashtags, Networks and Contagion.” *New Formations* 78(1):46–64. doi: 10.3898/NEWF.78.02.2013.

Smith, Jeffrey A., Miller McPherson, and Lynn Smith-Lovin. 2014. “Social Distance in the United States: Sex, Race, Religion, Age, and Education Homophily among Confidants, 1985 to 2004.” *American Sociological Review* 79(3):432–56. doi: 10/f6dxxp.

Solomon, Akiba. 2014. “Get on the Bus: Inside the Black Life Matters ‘Freedom Ride’ to Ferguson.” *Colorlines*. Retrieved October 26, 2021 (https://www.colorlines.com/articles/get-bus-inside-black-life-matters-freedom-ride-ferguson).

State, Bogdan, and Lada Adamic. 2015. “The Diffusion of Support in an Online Social Movement: Evidence from the Adoption of Equal-Sign Profile Pictures.” Pp. 1741–50 in *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing*, *CSCW ’15*. New York, NY, USA: Association for Computing Machinery.

Taylor, Derrick Bryson. 2021. "George Floyd Protests: A Timeline." *The New York Times*, November 5.

Tillery, Alvin B. 2019. "What Kind of Movement Is Black Lives Matter? The View from Twitter." *The Journal of Race, Ethnicity, and Politics* 4(2):297–323. doi: 10.1017/rep.2019.17.

Tufekci, Zeynep. 2014. "Social Movements and Governments in the Digital Age: Evaluating a Complex Landscape." *Journal of International Affairs* 68(1):1–18.

Tufekci, Zeynep, and Christopher Wilson. 2012. "Social Media and the Decision to Participate in Political Protest: Observations From Tahrir Square." *Journal of Communication* 62(2):363–79. doi: 10/f3xgtm.

Twitter, Inc. n.d. "About Twitter." *Twitter*. Retrieved November 23, 2021 (https://about.twitter.com/en.html).

Walgrave, Stefaan, Dieter Rucht, and Peter Aelst. 2010. "New Activists or Old Leftists?: The Demographics of Protesters." Pp. 78–97 in *The World Says No to War: Demonstrations against the War on Iraq, Social Movements, Protest, and Contention*.

Wang, Xin, Antonio D. Sirianni, Shaoting Tang, Zhiming Zheng, and Feng Fu. 2020. "Public Discourse and Social Network Echo Chambers Driven by Socio-Cognitive Biases." *Physical Review X* 10(4):041042. doi: 10.1103/PhysRevX.10.041042.

Welles, B. F., and Sarah J. Jackson. 2019. "The Battle for #Baltimore: Networked Counterpublics and the Contested Framing of Urban Unrest." *International Journal of Communication* 13:1699–1719.

Weng, Lilian, Filippo Menczer, and Yong-Yeol Ahn. 2013. "Virality Prediction and Community Structure in Social Networks." *Scientific Reports* 3(1):2522. doi: 10.1038/srep02522.

Wilkins, Denise J., Andrew G. Livingstone, and Mark Levine. 2019. "Whose Tweets? The Rhetorical Functions of Social Media Use in Developing the Black Lives Matter Movement." *British Journal of Social Psychology* 58(4):786–805. doi: 10.1111/bjso.12318.

Wu, Henry H., Ryan J. Gallagher, Thayer Alshaabi, J. L. Adams, J. R. Minot, M. V. Arnold, B. F. Welles, Randall Harp, P. Dodds, and C. Danforth. 2021. "Say Their Names: Resurgence in the Collective Attention toward Black Victims of Fatal Police Violence Following the Death of George Floyd."

Yang, Guobin. 2016. "Narrative Agency in Hashtag Activism: The Case of #BlackLivesMatter." *Media and Communication* 4(4):13. doi: 10/gf3gst.

Zappavigna, Michele. 2021. "Searchable Talk: The Linguistic Functions of Hashtags in Tweets about Schapelle Corby." *Global Media Journal: Australian Edition* 9(1).

Zitner, Aaron, and Anthony DeBarros. 2018. "The New Divide in Politics: Education." *Wall Street Journal*, November 10, 2.

Zulli, Diana. 2020. "Evaluating Hashtag Activism: Examining the Theoretical Challenges and Opportunities of #BlackLivesMatter." *Participations: Journal of Audience and Reception Studies* 17(1):197–216.

APPENDIX A: COMPARING GEOLOCATED ADOPTERS TO ALL ADOPTERS

Before I examined user behavior by geography, I needed to compare geolocated adopters to the overall adopter population, in order to see how representative my results are of the whole. The results of that investigation are described in this appendix. I find that geolocated adopters tend to be older and more connected on Twitter than the overall adopter population, many of whom come from outside the United States.

An examination of geotagging behavior by account birthdate demonstrates that geolocating was much more common among the first users to join Twitter; after this early period, geotagging behavior declines linearly as account lifespan becomes shorter. Thus, 3.04% of adopters who joined Twitter in 2010 have geotags, while only 0.51% of those joined Twitter in 2019 have geotags. The median birthdate for a geolocated adopter is March 2012, while the median birthdate for adopters overall was October 2015. This contrast may partially be explained by Twitter's interface. The geotagging feature is not obvious, and thus may require some investigation for users to turn on; however, once it is on, it may go unnoticed. The longer you are on Twitter, the more opportunities you have to turn on the feature. The median number of followers for a geolocated adopter was roughly double the median for all adopters (244 vs 460) and the median number of users followed by a geolocated adopter was slightly less than double (366 vs 597). Thus, geolocated adopters most closely resemble the Twitter "power-users" who were the earliest adopters of #BlackLivesMatter.

Geolocated adopters saw an earlier rise in the number of adoptions compared to overall adopters for the first two days of activity, before the overall adopter rate surged; for the rest of the study period, the cumulative proportion of adopters for each population were closely aligned. By midday on May 27th, roughly 36 hours into the diffusion event, 7.7% of geolocated adopters had already adopted, while 4.3% of adopters overall had. By the start of May 29th, this flipped, and 27.5% of overall adopters had adopted, while 23.4% of geolocated adopters had. A sharp increase in the number of overall adopters had not been matched by an equal increase in the number of geolocated adopters. It is not until May 31st that the geolocated adopters' cumulative adoption proportion matches that of overall adopters; from May 31st to June 2nd, the two curves stay relatively aligned.

The rise in adoptions by less-active Twitter users and presence of adopters from foreign countries in the overall adopter population can explain why the cumulative adoption rate varies for geolocated adopters. First, geolocated adopters represent a more active population of Twitter users. Given that May 26th through the 28th sees a trend towards adoptions by less-active users, who are less likely to be geolocated, this increase is less likely to be reflected in the geolocated adopter rate. Additionally, May 28th is the day in which #BlackLivesMatter begins to be widely adopted by users beyond the United States, as shown in Figure 17. The diurnal rhythm of human behavior means that the six-hours after midnight tend to show lower adoption rates among geolocated adopters in the United States, leading to a step-like rate of cumulative adoptions. The Overall Adopter rate lacks this pattern, in part because of the presence of adopters outside the United States with different diurnal patterns within the overall population. The sharp rise in foreign adoptions on the 28th masks the diurnal patterns of adopters within the United States; thus, the overall cumulative adoption rate increases overnight on the 28th just as the geolocated cumulative adoption rate stagnates. It is uncertain how many non-geotagged tweets came outside of the United States, however, a narrow majority of geotagged tweets during the study period came from outside the United States. Of 2.60 million geotagged tweets, 1.36 million (51.4%) geotagged tweets came from outside the United States. The top foreign countries represented in geotagged tweets were Brazil (300,969 tweets), Great Britain (223,445 tweets), Canada (74,791 tweets), Nigeria (64,732 tweets), and South Africa (59,832 tweets). The large proportion of geotagged adopters from outside the United States suggests that foreign adopters had a significant impact on the overall adoption rate.

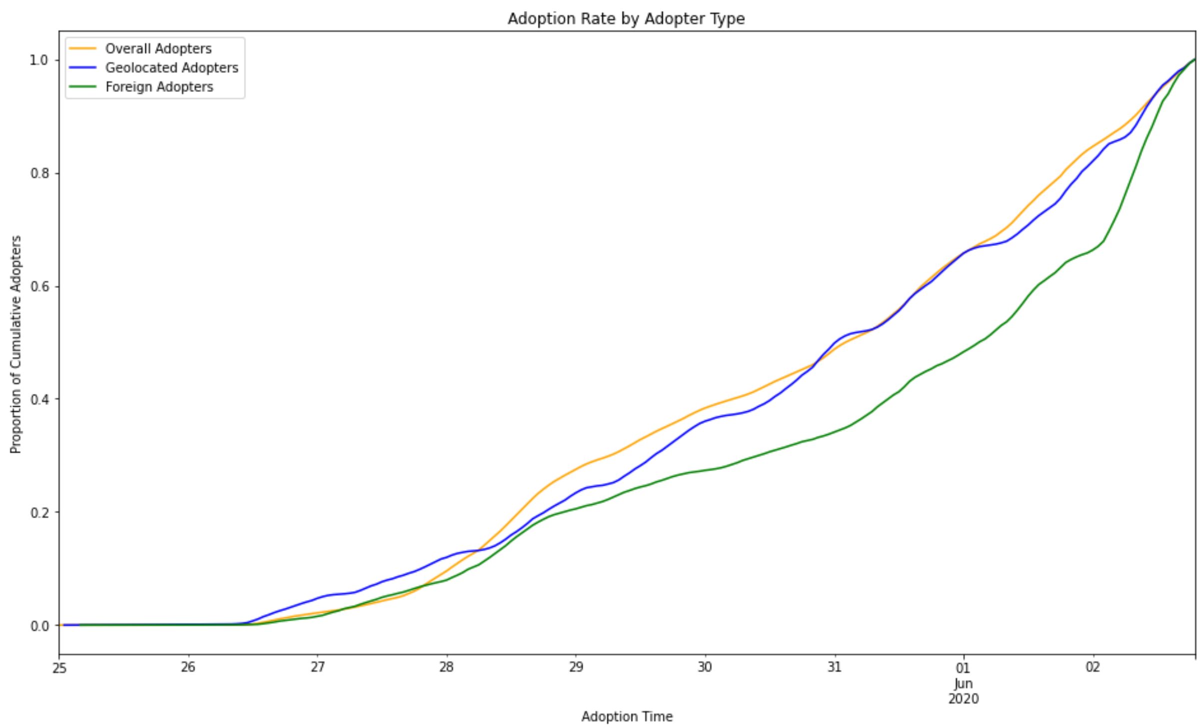

Figure 17: The cumulative adoption rate for geolocated adopters, adopters outside the United States, and the overall adoption population (as proportion of total adopters by June 2nd).

Now that I have contextualized how the geolocated adoption rate varies from the overall adoption rate, I note why the geolocated adopter population is still important for understanding diffusion: First, geolocated adopters' adoption patterns are not drastically different from those of the overall population. Despite differences between geolocated and other adopters' characteristics, the geolocated cumulative adoption rate never differed by more than twelve hours from the overall adoption rate. Second, while the role of adopters outside the United States is important, focusing on geolocated adopters may be valuable given that adopters from outside the United States are responding to events within the country. Finally, given that geolocated adopters are more active users who may be involved in seeding outbreaks, understanding their contagion dynamics is revealing of the contagion dynamics overall.

APPENDIX B: OBSERVATIONAL CONTAGION METRICS

Before developing their probabilistic contagion simulations, Fink et al. (2015) developed four network-based metrics to identify simple and complex contagions. Before conducting my simulations, I computed and analyzed these metrics. While I found inconclusive signs of complex contagion, this finding was contradicted by my later results from my simulations that homophily and external factors better explained hashtag diffusion over the entire study period. I have included the results of the original metrics, however, to be thorough, demonstrate the limitations of observational metrics in analyzing contagion, and to illustrate other interesting network properties. Specifically, these metrics demonstrate that hashtag diffusion was a rapid event, quickly permeating a large network of Twitter users, and reinforce the notion that BigCities individuals drove the online conversation.

Reinforcement Ratio

Fink et al. (2015) designed the reinforcement ratio to estimate the level of social influence a hashtag needed to be adopted. A higher reinforcement ratio is supposed to suggest that more social reinforcement is necessary before hashtag adoption. Fink et. al (2015) defined the reinforcement ratio as the ratio of adopters who followed at least two accounts who had already adopted at the time of their adoption; here I have modified this threshold to four, because almost as soon as data is available for each cluster on May 26th, nearly all adopters had at least two peer adopters. Even with this adjustment, the reinforcement ratio reached nearly 1.0 by May 27th for all clusters, as demonstrated in Figure 18. Nearly all new adopters beyond May 27th already had at least four or more neighbor adopters. Ultimately, I deemed the reinforcement ratio to be a poor metric for understanding contagion, because the sheer scope of #BlackLivesMatter's diffusion may mean that any sufficiently connected user would have followed at least four adopters. Nevertheless, I find this to be a useful metric, because it demonstrates the speed with which #BlackLivesMatter spread. Without also collecting information on the reinforcement ratio of individuals who did not adopt the hashtag, it is hard to say if the entire Twitter network was as saturated with #BlackLivesMatter messaging as the network of geolocated hashtag adopters, but if the overall adopter population saw reinforcement even moderately like the geolocated adopter population, this suggests a diffusion event which penetrated millions of users' networks within two days.

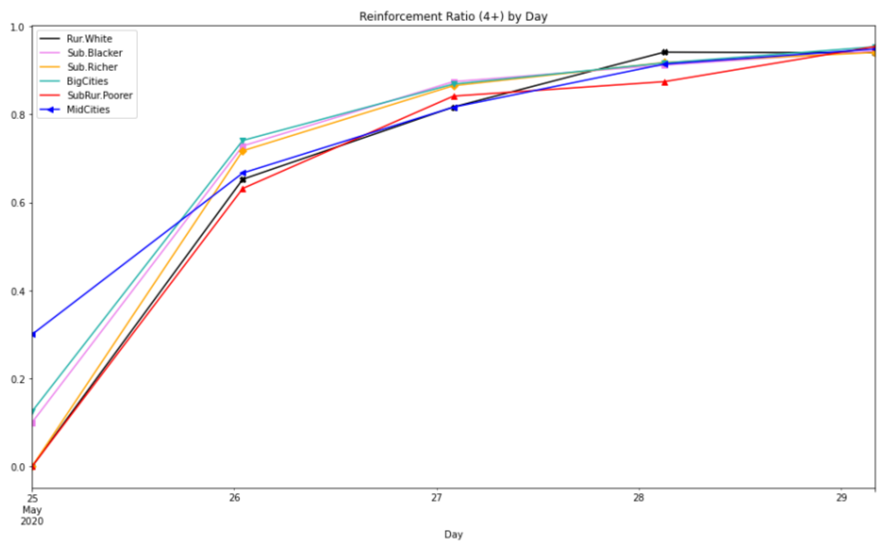

Figure 18: Reinforcement ratio (proportion of new adopters with four or more neighbor adopters at time of adoption) by cluster.

Within-Group Tie Ratios

Next, I examined the within-group tie ratio. The within-group tie ratio among the first m adopters is mathematically defined as:

$$\text{TR}\ (m) = \frac{1}{m(m-1)} \sum_{i \in \mathcal{S}_m} \sum_{j \in \mathcal{S}_m \setminus i} e(i,j)$$

where *e(i, j)* equals 1 if user *i* follows user *j*, and 0 otherwise. In simple terms, the within-group tie ratio represents the total number of following-relationships within a network, normalized by the total number of following-relationships that could exist in a perfectly-connected network — a network in which every user followed every other user. By examining tie-ratio among the first *m* adopters, we can determine whether a contagion began in a dense region of the network, which is a prerequisite to complex contagion. I note that whereas Fink et al. (2015) used this metric among all of the firs*t m* adopters of a hashtag, I use within-group tie ratio on subsets of adopters — specifically, those geolocated adopters in each cluster. This means that the generated tie ratios are most valuable in comparison over time, rather than to demonstrate the true density of the network of hashtag adopters.

I find that while the MidCities cluster had a higher initial density, the BigCities cluster's density remained elevated throughout the outbreak. Among the first one-hundred adopters, the MidCities cluster had the highest a within-group tie ratio; however, as shown in Figure 19, as the number of adopters approached five-hundred, the MidCities cluster's density sharply dropped while the BigCities cluster's density remained elevated throughout the study period. This finding has two implications: first, it suggests that MidCities adoption was initially a concentrated, local phenomena, with many users who already existed in the same phenomena. As more MidCities individuals adopted the hashtag, the emergence of adopters from MidCities beyond Minneapolis may dilute the strong interconnected effect seen among the earliest MidCities adopters. The second finding is related to contagion dynamics. Dense networks are conducive to the spread of complex contagion; thus, the fact that within-tie ratio's relative ordering of density nearly matched the ordering of per capita hashtag adoption beyond five-hundred adopters could suggest that social reinforcement is facilitating the higher levels of adoption of certain clusters. I interpret this finding in light of the results of my simulations, which suggest that homophily and external factors are what is driving hashtag adoption after the first forty-eight hours of adoption. These findings about within-group tie-ratio, then, lose much of their significance regarding evidence of contagion, but still reinforce the notion that the BigCities cluster dominated the #BlackLivesMatter geolocated network, as its high level of density would have facilitated fast information-sharing between BigCities adopters.

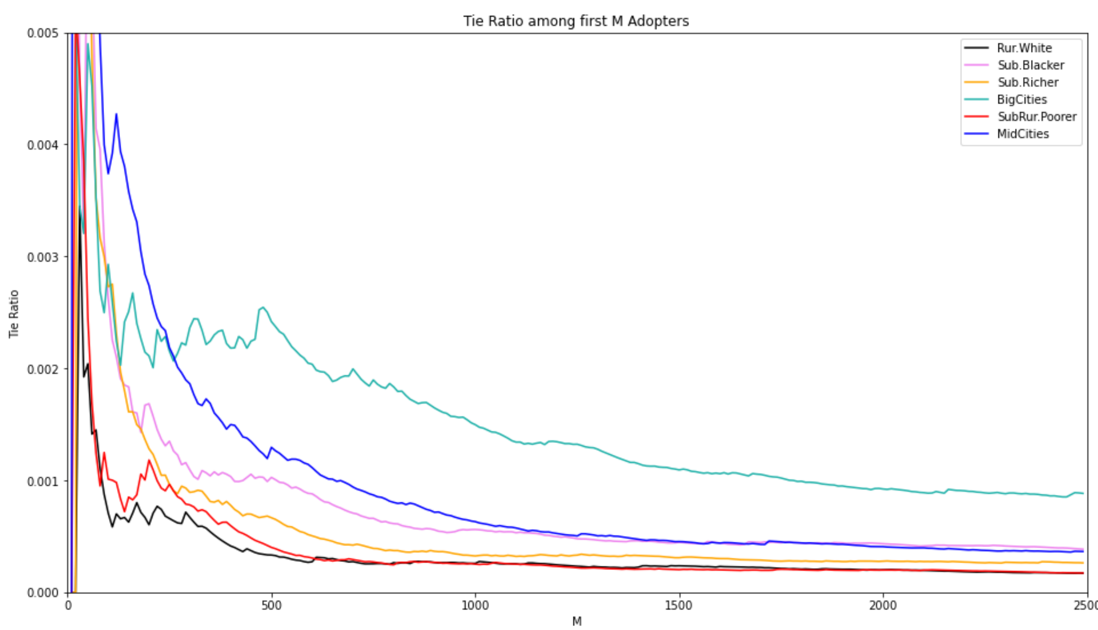

Figure 19: Tie-ratio as calculated among first M adopters within each cluster.

Cross-Group Tie Ratios

Next, I examined the tie-ratio across clusters, and I found that most clusters were highly connected to the BigCities cluster. I had initially expected that each cluster of adopters would be relatively isolated, with a few central nodes linking each cluster, but instead I found that each cluster's adopters had many connections to adopters in other clusters, which can help explain my finding in the previous chapter that each cluster had a relatively similar rate of proportional adoptions. Table 12 shows the tie-ratio for each cluster, normalized by the total possible number of ties which could exist per cluster. Not surprisingly, all clusters have high within-group tie ratios. The high tie ratio between all clusters and the BigCities cluster reinforces the notion that most adopters recieved much of their information on Twitter from the BigCities cluster. For example, adopters from the Sub.Richer cluster were even more connected to adopters from the BigCities cluster than they are to other Sub.Richer adopters. The high level of connection between clusters and from BigCities adopters to other adopters can further explain why different clusters had similar rates of proportional adopters over time -- adoption within BigCities may have led to similar rates of adoption in other clusters.

To illustrate the dominance of BigCities adopters, I constructed a visualization of the network of geolocated adopters, connected by follower-relationships. In Figure 20, each edge is colored by the cluster of the user being followed. This graph was constructed using the Force Atlas algorithm with Dissuade Hubs set to true, which pushes influential nodes towards the perimeter. The dominance of BigCities adopters in the online conversation is evident based on the large proportion of hubs colored turquoise (which represents a BigCities adopter). Note that the color coding for each cluster matches the colors in Table 12. Thus, while many users were adopting #BlackLivesMatter due to factors external to Twitter, the online Twitter conversation was dominated by users from large cities and urban areas.

Source → Target ↓	Rur.White	Sub.Blacker	Sub.Richer	BigCities	SubRur.Poorer	MidCities
Rur.White	1.41E-04	1.01E-04	1.26E-04	1.64E-04	1.06E-04	1.20E-04
Sub.Blacker	8.74E-05	2.64E-04	1.04E-04	1.64E-04	1.02E-04	1.19E-04
Sub.Richer	1.07E-04	1.03E-04	1.84E-04	2.13E-04	9.33E-05	1.27E-04
BigCities	8.73E-05	1.01E-04	1.44E-04	3.86E-04	7.96E-05	1.22E-04
SubRur.Poorer	1.02E-04	1.05E-04	1.01E-04	1.48E-04	1.58E-04	1.16E-04
MidCities	8.63E-05	9.88E-05	1.12E-04	1.66E-04	9.17E-05	1.87E-04

Table 12: Tie ratio for adopters in source cluster following adopters in target cluster. Highest tie-ratio for reach cluster highlighted dark gray, second-highest highlighted light gray.

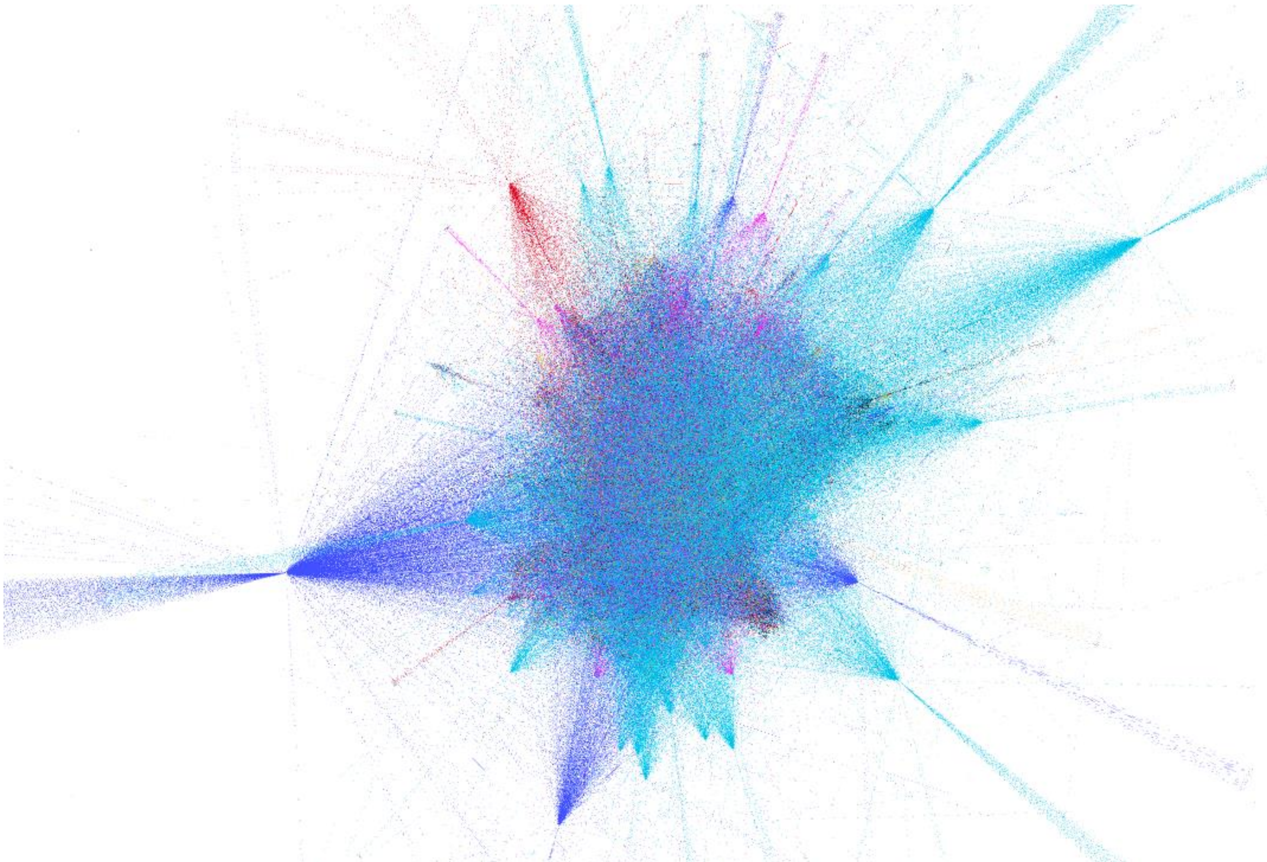

Figure 20: Geolocated adopters network constructed in GEPHI. Edges represent follow-relationships colored by source's cluster. Using Force Atlas 2.0 with scaling 2.0, gravity 1.0 and distributing attraction along outbound edges to push influential nodes towards perimeter.

Mean Overlap

The final metric defined by Fink et al. (2015) is mean overlap, which was the most important metric for determining contagion from observational data. The mean overlap score had inconclusive results regarding what kind of contagion dynamics were at play during hashtag spread, but I still include these results for thoroughness. The mean overlap metric was developed by Fink et al. (2015) and inspired by Barash (2011) and Weng, Menczer, and Ahn (2013)'s parallel findings that complex contagions tend to become structurally trapped inside local communities before reaching critical mass. Like within-group tie ratio, it helps measure network density; however, instead of measuring the density of ties among those who have already adopted, it looks at the density of ties among the neighbor-adopters for each adopter at a given time. To calculate mean overlap, for each first-time adopter in a given time period, you calculate the within-group tie ratio among just the first-time adopter's neighbors who have already adopted, and then calculate the mean of all such tie ratios.

Mean overlap is expected to drop rapidly following a period with a low adoption rate and during a period with a rapidly increasing adoption rate, in scenarios in which a contagion is complex. I omit the mean overlap for each cluster given that each cluster's result was similar to the result for the overall population of geolocated adopters, which is displayed in Figure 21. While it is true for each cluster that overlap drops over time and adoption rate increases over time, there is not a distinctive, perfectly timed relationship between these variables, as Fink et al. (2015) found when they used this metric for the hashtag #BringBackOurGirls. For example, the first sharp drop in mean overlap is at the beginning of the day May 27th, but the following twelve hours of contagion see roughly the same rate of new adoptions as the prior twelve hours.

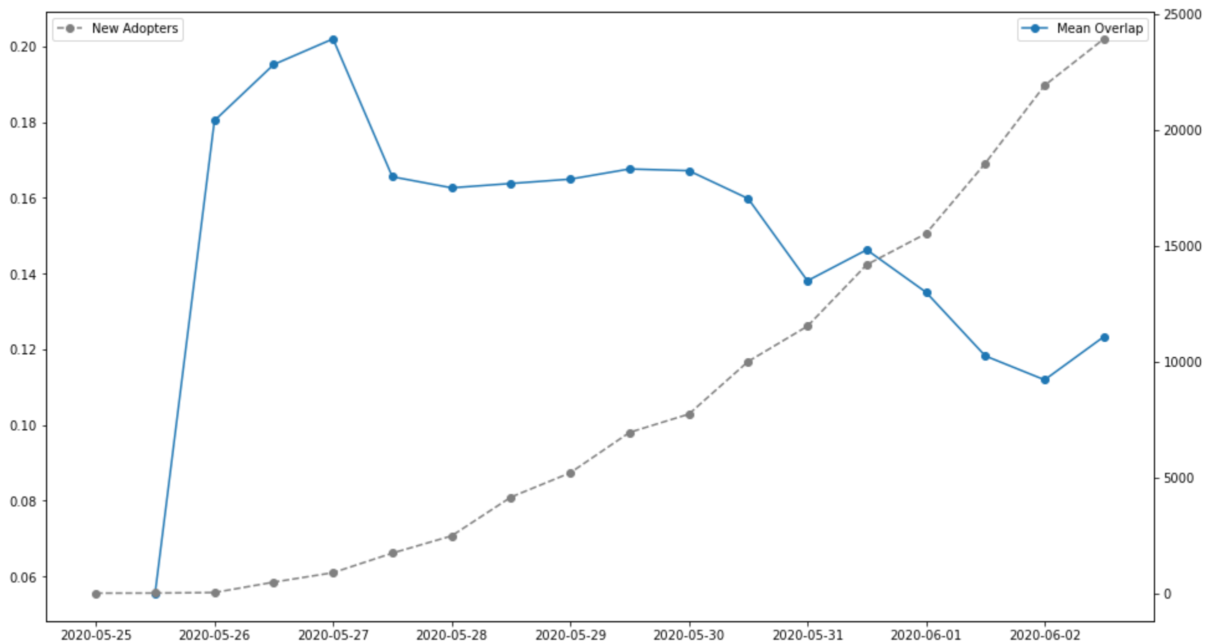

Figure 21: Mean overlap for all geolocated adopters plotted alongside cumulative adoptions.

In light of my finding that external factors best explain hashtag adoption, I argue that the mean overlap does not show signs of critical mass. Instead of slowing and then sharply increasing, the rate of new adoptions continues to increase throughout the study period, as adoption of #BlackLivesMatter transitions from being a Twitter-specific diffusion event to a much larger, online and offline event.

Summarizing Findings

While the metrics defined by Fink et al. (2015) may not reveal the contagion dynamics that I sought to study, they do reveal important details of the #BlackLivesMatter geolocated network. Specifically, they reveal the speed with which #BlackLivesMatter overtook Twitter, and reinforce the notion that #BlackLivesMatter usage was predominantly an urban activity. Future researchers should use caution when attempting to use these metrics, as they may falsely reveal signs of contagion when none exists, but these metrics can be useful for illustrating nuances in network features.

APPENDIX C: FULL TRAINING RESULTS

Cluster	Reference Model	Simple Model Parameters, Results	Complex Model Parameters, Results
Rur. White	{'q': 8e-06} 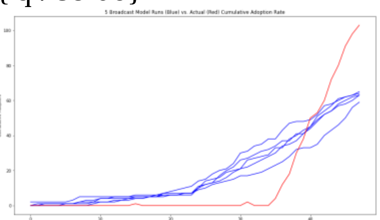 Mean (over 5) Cost: 263.7	{'p': 0.003, 'q': 8e-06, 's': 0.2} 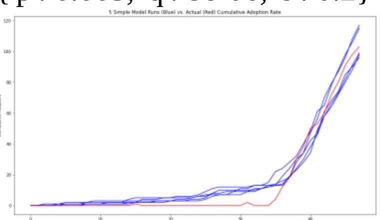 Mean (over 5) Cost: 39.6	{'k_0': 100.0, 'p_hi': 0.38, 'g': 0.15, 'q': 8e-06, 's': 0.15} 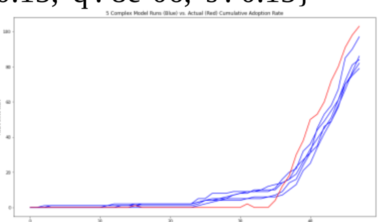 Mean (over 5) Cost: 71.9
Sub. Blacker	{'q': 1.2e-05} 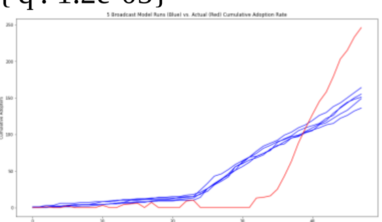 Mean (over 5) Cost: 1479.5	{'p': 0.005, 'q': 1.2e-05, 's': 0.1} 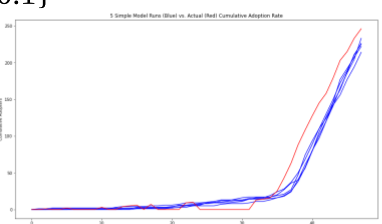 Mean (over 5) Cost: 338.8	{'k_0': 50.0, 'p_hi': 0.2, 'g': 0.15, 'q': 1.2e-05, 's': 0.1} 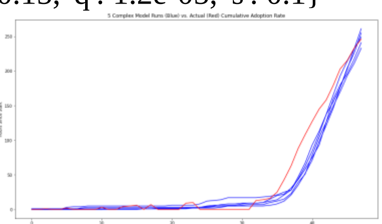 Mean (over 5) Cost: 298.7
Sub. Richer	{'q': 1e-05} 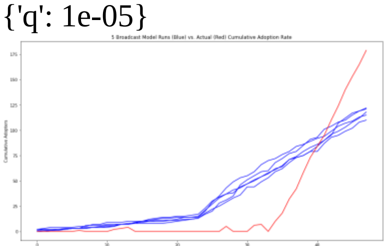 Mean (over 5) Cost: 791.7	{'p': 0.0038, 'q': 1e-05, 's': 0.1} 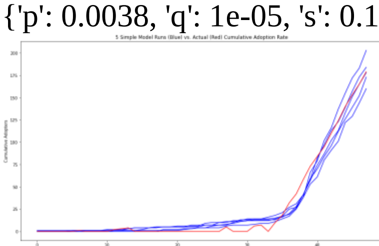 Mean (over 5) Cost: 67.9	{'k_0': 80.0, 'p_hi': 0.34, 'g': 0.15, 'q': 1e-05, 's': 0.2} 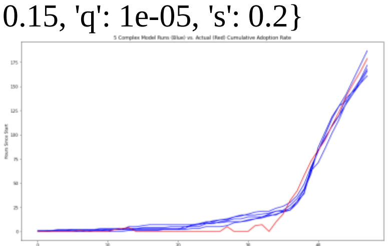 Mean (over 5) Cost: 56.2
Big Cities	{'q': 1e-05}	{'p': 0.0046, 'q': 1e-05, 's': 0.1}	{'k_0': 60.0, 'p_hi': 0.14, 'g': 0.15, 'q': 1e-05, 's': 0.2}

	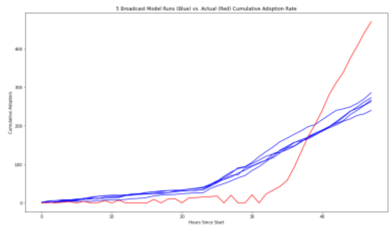 Mean (over 5) Cost: 5059.5	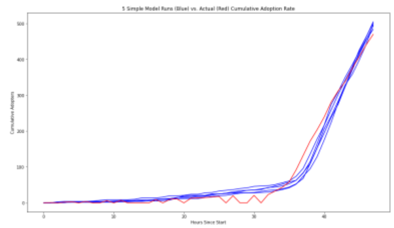 Mean (over 5) Cost: 448.7	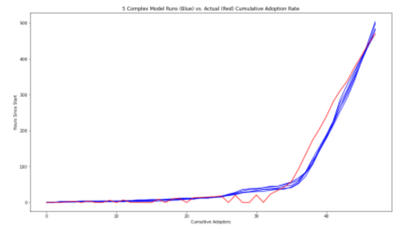 Mean (over 5) Cost: 555.1
Sub Rur. Poorer	{'q': 8e-06} 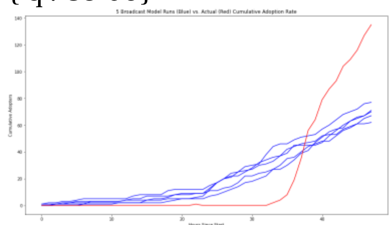 Mean (over 5) Cost: 562.7	{'p': 0.0046, 'q': 8e-06, 's': 0.1} 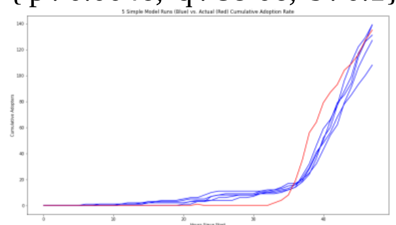 Mean (over 5) Cost: 121.1	{'k_0': 50.0, 'p_hi': 0.2, 'g': 0.15, 'q': 8e-06, 's': 0.1} 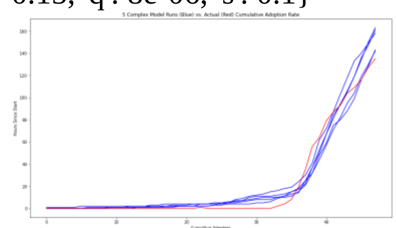 Mean (over 5) Cost: 67.2
Mid Cities	{'q': 9e-06} 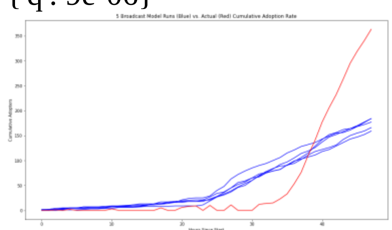 Mean (over 5) Cost: 3765.0	{'p': 0.0046, 'q': 9e-06, 's': 0.15} 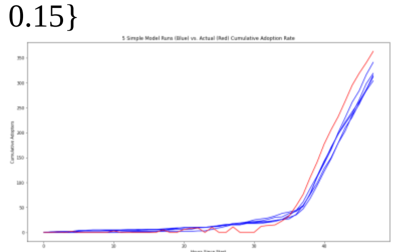 Mean (over 5) Cost: 497.5	{'k_0': 50.0, 'p_hi': 0.22, 'g': 0.15, 'q': 9e-06, 's': 0.1} 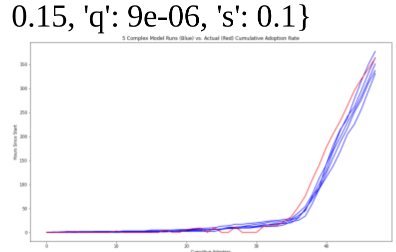 Mean (over 5) Cost: 378.8

Table 13: The results of training each diffusion model on each cluster of geolocated adopters.

The trained parameters and mean least-squares cost are included for each.