Title
The People Formerly Known as the Audience: The Political Effects of Crowdsourcing News about Violence

Author
WIRTSCHAFTER, VALERIE

Publication Date
2021

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA

Los Angeles

The People Formerly Known as the Audience: The Political Effects of Crowdsourcing News about Violence

A dissertation submitted in partial satisfaction

of the requirements for the degree

Doctor of Philosophy in Political Science

by

Valerie Wirtschafter

2021

ABSTRACT OF THE DISSERTATION

The People Formerly Known as the Audience:

The Political Effects of Crowdsourcing News about Violence

by

Valerie Wirtschafter

Doctor of Philosophy in Political Science

University of California, Los Angeles, 2021

Professor Barbara Geddes, Chair

How does new information affect citizens' political attitudes and personal beliefs? Past research has explored this question with mixed results: while some suggest that information can alter beliefs and policy positions, others argue that political opinions are stubbornly resistant to change, regardless of the potential value of new information. Building on this expansive body of research, I argue that both individual-level uncertainty and the information medium play a critical role in shaping when information matters. While I expect citizens who are more uncertain to be more likely to reevaluate their beliefs in the face of new information, I argue that in specific contexts, this is particularly true if that information comes from a source that is viewed as trustworthy, like their community.

Using data from two survey experiments and a large-N original dataset of Facebook community pages, my dissertation explores this dynamic in Brazil with information about violence, a persistent development challenge. Conditional on respondents uncertainty about community violence, I find that community-based information shared through social media

has significant and countervailing effects: (1) it significantly decreases anxiety about crime and support for the police among more "uncertain" individuals, but increases it for less "uncertain" individuals. I find no such heterogenous effects among respondents who receive information about violence in their community from government sources. Examining the divergence in high and low certainty individuals who receive the social media information is beyond the scope of this dissertation, but what is clear from these findings is that the source of the information plays a central role in shaping whether people are willing to update their beliefs.

I argue that this observed shift in beliefs is due to respondents' willingness to trust information that comes from "people like them" over more traditional purveyors of information, such as their government. I test this through a conjoint experiment, which randomly varies features of information about violence and government services and asks participants to select which they would trust more. Across both types of information, participants are significantly less likely to trust information that comes from a government source, and more willing to believe information that has a wide reach on social media and that is shared in real time. Uncertainty appears to play little role in respondents' weighting of these various attributes.

In the final chapter of my dissertation, I investigate how the transition from traditional to social media sources for information shapes individuals' beliefs about violence in their communities, which in turn shape their political preferences. Although respondents appear more willing to update their beliefs based on information from social media, that information does not accurately reflect real world instances in violence. An analysis of content from community level Facebook groups reveals clear biases in the content shared on these platforms. Social media users based in areas that are neither very violent nor very peaceful are more likely to share content about violence. This finding suggests that those who have less certain information about particular aspects of the real world will be more likely to increase the inaccuracy of their understanding if they rely mostly on social media for information.

This research not only complicates existing theories about information processing, but also explores these dynamics in an understudied and rapidly transforming political environment where social media usage is rapidly accelerating.

The dissertation of Valerie Wirtschafter is approved.

Graeme Douglas Blair

Chad J. Hazlett

Daniel N. Posner

Zachary Steinert-Threlkeld

Barbara Geddes, Committee Chair

University of California, Los Angeles

2021

To my family, without their love and support,

this dissertation would not be possible.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGMENTS

Completing a Ph.D. takes a village, and there are countless people who helped me along the way. I want to thank my chair, Barbara Geddes, who took a chance on an intellectually curious, if naive student who had never read a political science article at the start of graduate school and molded me into the researcher I am today. Thank you for all your handwritten line edits, words of encouragement, and chats about political science and life. Barbara challenged me to think deeper about the intellectual ties that anchor this research to political science, while encouraging me to pursue a topic that is wholly and unequivocally my own. She surely knows this by now, but Barbara's example is truly a model to follow, and I am so happy to have had the opportunity to learn from your guidance, brilliance, and – most importantly – kindness. I kept me moving forward throughout the program, even at times when I wanted to call it quits, and I am forever grateful.

My dissertation also benefitted from a stellar committee, who provided valuable insights throughout the duration of this project and played a critical role in shaping my time at UCLA. Chad Hazlett, thank you for showing me early on that I *am* a capable quantitative researcher, despite every part of me saying that was wholeheartedly not the case! As in-demand of a mentor as he is, thank you for making time to support this project and more importantly to support my development as a researcher into avenues that I previously felt were impossible. Dan Posner, thank you for your words of encouragement, thoughtful prodding, and methodological support for this project. Thank you for introducing me to the cutting edge in political science research and urging me to chart a similar course. Zachary Steinert-Threlkeld, thank you for you for engaging so thoughtfully with this research project, and for serving as a substantive support system for a topic that is so new to the political science canon. Graeme Blair, thank you for always challenging me to improve aspects of the research design, and for providing a model to help make my work more sophisticated, rigorous, and thorough. I am a better researcher as a result.

Other faculty members who did not serve on my committee also provided significant

guidance and mentorship along the way. Thank you to Darin Christensen, for your assistance in thinking about the framing for this project early on, and Rob Trager, for your support throughout the program and the opportunity to work on a range of other exciting project. Outside of UCLA, I wanted to thank Natalia Bueno, who has been a fountain of guidance and wisdom, and who I only wish I had met earlier on so that she too could have joined my committee; Julia Sweig, who has been my number one champion since my Council on Foreign Relations days; and Chris Meserole, who has been nothing but supportive of my dissertation and played a vital role in helping me find an intellectual home at the Brookings Institution.

Although my time in the field was cut short by the coronavirus pandemic, this research also benefitted greatly from the insights of those who graciously welcomed me into their lives in Brazil. I want to thank Daniel Elder, who first introduced me to the idea of the crime crowdsourcing apps in Rio; Maria Isabel MacDowell Couto, Guilherme Pimentel, Marquinhos Vinícius, who spoke to me candidly about how they work in practice; Pablo Nunes, whose vital work on community Facebook pages inspired the fifth chapter of this dissertation; Cesar Zucco and Elizabeth Stein, who brought me into their UCLA community and opened doors for me in the field; and finally, Ana Pellegrino, Daniel Cerqueira, Andre Batista, Julia Guerra, and Joana Monteiro who shared vital insights with me on the challenges of public safety and citizen security in Rio de Janeiro, Brazil. Jessie Bullock, who was also conducting field work at the same time as my first trip to Rio, offered valuable insights on working in the field and was a source of fun and friendship while abroad. Finally, I want to thank Quaest Pesquisa e Consultoria, and specifically Felipe Nunes and Renata Salvo, who helped me think through survey design and expertly fielded the survey during a once-in-a-generation pandemic.

I would also be remiss if I didn't thank my Brazil "family," who welcomed me with open arms every time I returned to Rio and always made me feel like I had a home. Alex Menezes, thank you as always for your mentorship, advice, and excellent dining taste. Paola de Salvo, your energy and friendship were exactly what I needed while traveling alone. Thank you for

always showing me a great time and for being an incredible source of support and insight. Gregório de Matos, thank you for being such a great friend and for your keen interest in my research, even when I worried that no one would care about what I had to say.

Graduate school can be an isolating place – and a dissertation project even more so – but friends both at and outside of UCLA helped to provide levity and guidance throughout. Tanu Kumar, you are equal parts role model and best friend. We've grown together since the first day of college, and I am so grateful for your friendship, advice, and support. Cesar Martinez, you are my rock and partner in crime. Thank you for reading draft after draft, talking through research ideas, and truly being the kindest and most supportive friend. Venting to these two became an at least weekly (sometimes daily) habit, and I wouldn't have made it through the program without them. Aaron Rudkin, thank you for being such an outstanding mentor, understanding teacher, and wonderful friend. I am forever indebted to you for every ounce of support (both moral and intellectual) you provided. Fernando Mello, your support since day one has been invaluable. Thank you for your guidance, suggestions and friendship; and for so graciously translating my survey.

I also wanted to thank other friends at UCLA and in Los Angeles who made the Ph.D. process more enjoyable and whose guidance and support I benefitted from along the way: Jenny Hamilton, Brigid Morris, Frank Wyer, Tamar Zeilberger, Mariana Barragan, and Carson Bruno. To Linda Kinstler, who was my unrelenting champion from afar, and the first one to welcome me with open arms after my surprise move to Washington, DC. To "the Girls," Jeannine Donnelly, Malini Sen and Tom Wallack, whose cheering bookended my Ph.D. and presence kept me sane during the final sprint to the finish. I also want to thank Taryn Black Sarnoff, who likely won't think she deserves to be in here, but she does more so than she knows. Finally living in the same city as my best friend of 20 years and watching her little girls grow up was the biggest treat I could ask for and meant more toward getting me through this process than she will ever know.

I also want to thank my dear family, who, after nearly a decade on the East coast

prior to my Ph.D., I never thought I would live so close to again. My parents, Elise and Gary Wirtschafter, provided Sunday meals every week, sat through practice prospectus presentations, and celebrated every milestone, including becoming "Chicken Karma MA." Thank you for the value you have placed in my education and for the unconditional love that never went unnoticed. A special shout of to my mom, who taught (and is still teaching me) how to write, and who will likely find several grammatical errors in these acknowledgements. Thank you for reading and editing every academic document I've likely ever produced in my life. That is a gift that most people do not have, and I am forever grateful. To Eric Wirtschafter, you are the kindest most thoughtful big brother I could ask for and to live nearby you was an absolute joy. Gloriana Mohr, thank you for being the fearless leader of the household. I knew you from bopping in and out of LA for years, but to really get to know you during my dissertation was wonderful beyond expression. I love the little nuggets you and Eric brought into this world, Landon and Liam, and to get to be with them from the start is a dream that I never thought possible. To Carla Wirtschafter, my roommate, duplex mate, best friend, and sister. You've always been my number one champion, and with you by my side, I know that anything is possible, even if it seems daunting. To have a sister who is your other half is an inexplicable feeling and living side by side with her truly made me feel whole. Thank you also for bringing Joe Thorpe into my life. I couldn't imagine a better, weirder, more amazing man with whom to share the love of my big sister. A special shout out to Lady Beatrice and King Herbert too, the actual leaders of your household who gave the corgis lifelong pals, and fiercely guarded the house at all times.

Finally, I wanted to express my unending gratitude to Cody Giddings, who stood steadfast with me through it all. We first met during Math Camp, but little did I know then that while re-learning probabilities, I would find the most perfect partner to walk through the rest of life with too. I could not have finished this program or this dissertation without you. Thank you for listening to craziest ideas with patience and skepticism, for encouraging me to pursue avenues of inquiry I didn't think possible, for coaching me through countless tears,

and for motivating me to get this done, while simultaneous taking on other responsibilities. Thank you for being the best dad to our little corgis, Otis and Frankie, and the best partner to me. The Ph.D. process was so hard, but I have no regrets simply because it brought me to you.

VITA

2012 B.A. (History and Latin American Studies), Bowdoin College.

2018 M.A. (Political Science), University of California, Los Angeles.

PUBLICATIONS

Blair, Graeme, Darin Christensen, and **Valerie Wirtschafter**. "How Does Armed Conflict Shape Investment? Evidence from the Mining Sector." *Journal of Politics*, 2021.

TEACING

Instructor of Record PS139: International Politics in the Digital Age (2021)

1

Introduction

> "Tiros no Pavão-Pavãozinho, em Copacabana (05:32) #TirosRJ #FogoCruzadoRJ"

The alert confirmed my suspicions. Gunshots right outside my window from *Rua Saint Ramon* – one of the main entrances to the relatively peaceful favela of *Pavão-Pavãozinho* in Rio's south zone – had woken me up early in the morning. I had been in Rio barely a week for dissertation research, and alone in my apartment, I instinctively took to social media to see what was happening right outside. After *Fogo Cruzado* (Cross Fire, in English), a crowdsourcing application focused on violence across some of Brazil's most populous cities, sent out an alert warning of gunshots in the Copacabana neighborhood of *Pavão-Pavãozinho*, I no longer felt concerned. Searching online, I knew that I was not alone in hearing these shots and felt calmed knowing that others were around me, feeling the anguish that I felt, despite being by myself.

In subsequent conversations about my experience over the next few days, I quickly saw two perspectives emerging among the Brazilians I interviewed. According to government officials and police officers with whom I spoke, access to this type of information could only be harmful, causing mass hysteria among citizens in ways that were incompatible with government data on violence – and the reality as they saw it. By contrast, among citizens

and activists, this type of information was a potential life line from the community and for the community – a trusted source with no agenda but to keep citizens safe and out of the way of cross fire.

From these conversations, it became clear that this tension was part of a much larger debate taking place across the city of Rio de Janeiro, Brazil, and many other cities across the world: With the rise of the internet and social media platforms, large amounts of previously untapped data, independent of traditional information gatekeepers, are now available for citizens to incorporate into the formation of their attitudes and beliefs. Do citizens assign more value to this content than more traditional purveyors of information? How does this crowdsourced content shape their perceptions about their surroundings? And does this new source of information paint the same sorts of pictures as previous, more official channels of information, often provided by the government or a select few elites, or are there distortions that occur along the way? These are the broad questions that help to guide my research, and which I seek to answer in subsequent chapters of this dissertation.

To explore these questions, I focus on the municipality of Rio de Janeiro in Rio de Janeiro State, where social media networks have evolved at an unprecedented speed to fill information holes left by the state government. The most obvious example of this phenomenon has occurred with respect to information about violent crime. Since the World Cup in 2014 and the Olympics in 2016, applications and online resources have rapidly expanded the availability, frequency, and breadth of information about violent crime – a major challenge in Brazil's second largest city, which some journalists argue is "a city at war." Various applications allow users to report gunshots and police brutality. There are networks for monitoring police checkpoints and for sharing on-the-ground footage from police operations. There is at least one community Facebook page for nearly every neighborhood in the municipality. Unlike traditional media, these platforms and pages are citizen run and publicly accessible, providing a bottom-up information source that was previously unavailable. Although the use of crime-focused applications is the most established in Rio de Janeiro, they are be-

coming increasingly popular in other parts of the world. In the United States, for example, citizens can find information about crime in their neighborhood through the Ring Camera App, if they own the product, and through a crowd-sourcing app called Citizen. The use of neighborhood Facebook pages is less common in the U.S., but it is similar to the platform NextDoor, without the required address verification.

1.1 Existing Theories

The role of social media in shaping public opinion is an understudied topic in political science research.What research does exist largely draws on empirical evidence from a U.S.-based context, which is unique for a variety of reasons. Outside of the U.S. context research is still rarer, and to my knowledge broad theories of public opinion formation in the digital age are non-existent. Yet, existing literature on how persuasive communications shape political attitudes and actions, which straddles two separate fields in political science that rarely interact, can help inform analysis of these dynamics. The literature in American Politics, influenced by the field of psychology, evaluates how persuasive communications shape political attitudes, with limited differentiation by medium and with a focus on how elites shape mass opinion. The literature in Comparative Politics, largely shaped by development economics, focuses on how information provided to voters impacts accountability among politicians. This research often fails to explain how the prior beliefs of respondents affect their reception of the provided information.

The standard theory of public opinion formation in American politics argues that people update their beliefs when exposed to new information, but that the amount of updating depends on how firmly they hold their initial beliefs. Extensions of this basic theory argue that people filter this information based on its consistency with their general worldview and their trust in its source. However, these theories and the empirical evidence supporting them were developed before the era of social media, which has transformed the information

environment from a monologue — where elites speak to the average citizen — to a multi-directional conversation — where the average citizen can now easily speak back and amongst each other. In my dissertation, I investigate whether the differences between social media and more elite-driven forms of political communication have altered the process of opinion formation and opinion change, and thus political outcomes affected by that opinion, such as vote choices or support for specific types of policies.

Established literature in Comparative Politics theorizes that in electoral democracies more informed citizens will be better able to hold leaders accountable. Thus more information, should in theory, make citizens more informed and better equipped to sanction or reward elites. To date evidence in support of this theoretical expectation has been mixed, with more recent rigorous empirical work demonstrating no discernible effect of information on political behavior. Although in many of these studies, information tends to be source and context independent (and assumed to be true), political attitudes operate in very specific contexts, and information holds different power depending on where it originates. In my dissertation, I draw inspiration from this research by incorporating into the analysis the personal motivations that shape an individual's willingness to update based on new information. I also explore the context – and perceived credibility – behind how information is received.

1.2 My Argument

The theory that I set forth in **Chapter 3** is influenced by the first empirical chapter and applicable to the remaining two empirical chapters. I summarize each empirical chapter below in **Section 1.3**. I take as a starting point the literature on attitude formation, persuasion, and public opinion, which is well-established and heavily theorized in the American Politics field. Bayesian theory argues that new information, if it challenges prior beliefs and is salient to individuals, should lead to a reassessment of prior opinions, implying that new information should catalyze a shift in policy preferences. To be more literal, an individual may come to

a situation with a prior idea about the probability of observing some phenomenon in nature but will update their beliefs about that phenomenon based on new information that they receive.

In Bayesian theory, the relationship between opinion change and prior attitudes depends on the certainty of an individual's prior beliefs, which is mathematically expressed as the distribution of prior beliefs. If someone is highly certain about their prior beliefs, evidence inconsistent with that belief will produce little updating. However, if someone is less certain about their prior beliefs, evidence that challenges that belief may lead to a change in attitudes.

Weak priors (that is, uncertainty), however, will not automatically induce updating: the nature of the information also matters, a factor that largely goes unaddressed in the original Bayesian model of information processing. Some research, however, has already complicated Bayesian theory to incorporate features that make people more or less receptive to new information. For example, people are more likely to accept new information from those they recognize as experts (McGuire, 1985) or from co-partisan sources (Zaller, 1992). These theories of opinion change were developed to explain the effect of one-way communication from elites to ordinary people, and they lead to clear-cut expectations about how exposure to new information in the media influences the opinions of those exposed. When elites share particular values, the citizens who are most exposed to elite discourse are most likely to adopt the same values. When, however, elites disagree, then citizens tend to adopt the views of those they find more trustworthy (co-partisans in the U.S.) and reject the information provided by those they view as less trustworthy. Regardless of whether elites are unified or divided, this process tends to result in popular views biased in the same direction as those of elites.

Since the bulk of this research was conducted, technological change has introduced new factors that shape the information environment, which in certain contexts, might affect the receptiveness of individuals to new information. In particular, I argue that what is unique

about social media is that it has transformed the information environment from a one-way monologue delivered by elites and media to ordinary people into a multi-way conversation, where elites can speak directly to citizens, citizens can speak back to elites, and citizens can communicate amongst themselves directly. The gatekeepers to the flow and spread of information have largely been eliminated or reduced to mere actors alongside the general public.

The contemporary transition from elite-to-mass media flows to multi-directional social media raises an important question: how should we expect individuals to update their views in response to information received from social media? I argue that in young democracies, with less established governance structures and independent media, this broadening of the information environment can enhance the credibility of new information because people have more trust in information that comes from "people like me" than from elites. This indicates a greater likelihood that information from social media sources will lead to updating, particularly among uncertain individuals who might be more prone to changing their beliefs in response to relevant stimuli.

Although this enhanced credibility may increase the likelihood that new information leads to updating, updating may not bring opinions closer to accurate understandings of the world, in contrast to what Bayes assumed. Due to the way content is disseminated and policed, every day users will—unsurprisingly—share and boost the signal of content that is interesting to them and consistent with their worldview. This introduces biases in the information environment based on the preferences of the individuals sharing content.

As a result, the peer-to-peer nature of social media may be more trusted, but it will also be more prone to error or distortions shaped by the interests of those who disseminate it than information provided by conventional media following norms about objectivity and source checking. In particular, I expect users to be more inclined to share and disseminate content when it is shocking, but uncommon. This builds on research by Kahneman and Tversky (1979), who argue that due to cognitive biases, people will either ignore or more heavily

weight improbable events and exaggerate their importance far more than they would, if they thought in terms of probabilities. Thus, a single shocking incident may be overweighted simply because it is probabilistically unlikely. Though this reality is possible for shocking news on television or in newspapers, the bias may be even larger in social media than in traditional media due to the amateurization of content creators. This has significant implications for political behavior in communities where the majority of the population acquired their information primarily from social media sites.

1.3 Overview of the Dissertation

In this section, I provide an overview of the three chapters that constitute the empirical components of my dissertation. The first empirical chapter relies on cross-national survey research; the second chapter on an original survey, featuring two embedded survey experiments; and the third, on a large-N analysis of social media content collected from local Facebook pages.

Chapter 2: All the News That's Fit to Tweet

This chapter serves two purposes: (1) it provides background information on the transformation of communications technologies and why social media may have different consequences for opinion formation than past innovations; and (2) it provides suggestive empirical evidence to motivate a broader theory, summarized in the previous section, based on the idea that citizens may be more willing to trust information when it comes from a source viewed as "just like them" – a reality made possible in the digital age. In particular, I ask: (1) how does social media differ from previous kinds of communication flows and what might be the implications of this for politics? and (2) are those implications evenly distributed or are there areas where we might expect the difference between social media platforms and other communications methods to matter more for political behavior?

In the first part of the chapter, I describe the technological landscape, highlighting how

new communications technologies transform political processes as a result of shifts in the information environment. For example, with the printing press the speed with which information could be published was dramatically altered, allowing for a more rapid dissemination of information. The radio marked the first time that politicians could speak to the masses on a 24-hour-cycle, diminishing the value of intermediaries like media executives or journalists. These changes, not surprisingly, catalyzed equally dramatic political shifts, which are well documented in social science and historical literature (see for example, a variety of studies about the political effects of the radio (Adena et al., 2015; Rigterink and Schomerus, 2017), and cell phones (Guriev, Melnikov and Zhuravskaya (2019); Shapiro and Weidmann (2015))).

Despite major differences in technological capabilities, the flow of information from elites or experts to ordinary citizens predominated in all earlier forms of information transmission. With the growth of social media platforms, the direct line to the masses has shifted to a far more dynamic conversation, with the masses able to communicate back – and amongst each other – in unprecedented ways. I argue that this is the critical difference between social media and other forms of communications, and this difference could have significant implications depending on the context in which social media users operate. With respect to social media, however, social science research has to this point paid only limited attention to how these dynamics play out in the non-U.S. contexts, despite the fact that it is in countries in Asia, Africa and Latin America where their usage is rapidly growing and where citizens are more likely to use these platforms for news.

Current empirical trends suggest that ordinary people are more likely to use social media as their main source of news when trust in political leaders and mainstream media are low. I highlight several trends that distinguish environments where social media usage is growing fastest from those where the majority of empirical studies have occurred. These trends show that citizens are more likely to get their news from social media in countries that are less democratic, less developed, and have less press freedom. Using survey evidence from LAPOP, I find that in more developed countries such as the U.S. (and also Canada), citizens who

use social media for news are more likely to trust the media and the government, whereas in non-U.S. contexts, citizens who use social media for politics are significantly more interested in politics, and less likely to trust the media and the government.

Consequently, we should expect social media to play a larger role in shaping public opinion, and political outcomes affected by such opinion, in countries where people trust political leaders and established media less. The countries where research on public opinion has largely occurred, however, are more developed and democratic. In addition, thay have higher press freedom. In this dissertation, I focus on a country – Brazil – that is more representative of the range of levels of trust in media and political leaders. Brazil is an ideal place for this research because it is considered fairly average by all measures including those related to development, democracy, and press freedom.

Chapter 4: Trust Thy Neighbor: How Social Media Shapes Political Attitudes about Violence

In this chapter, I test the theoretical claims made in **Chapter 3** about how the uncertainty of previously held opinions affects the likelihood of opinion change after exposure to new information and how the credibility and trust-worthiness of information sources affect the likelihood that those exposed to new information will accept it. I conduct these tests in Brazil, where many people get most of their news from social media, and my tests focus on a subject that is highly salient to most Brazilians and a common focus of social media conversations and political debate: violent crime. I seek to answer three specific questions: (1) how much does the strength of previous beliefs affect opinion change? (2) are respondents more likely to change their beliefs about the level of violence if new information comes from people like them? and (3) if beliefs are more malleable in response to new information from social media than from other sources, what specific characteristics of social media make it more effective?

My expectations for this chapter are twofold: (1) we should expect new information to effect greater change on the attitudes of citizens who are less knowledgeable about the subject at hand and/or less certain about their knowledge than those who are more knowledgeable;

and (2) social media information, due to the nature of its sourcing, is more likely to lead to attitude change because its sources are considered more trustworthy.

To test this theory, I rely on two survey experiments. The first provides respondents with information from either social media sources or an official source (in this case, the government). Due to ethical concerns, I provide respondents with accurate information about their community and about the information generating process. As a result, the actual treatment is a bundled one, which varies timing and source, as well as the content of the information people receive. I also ask respondents a series of knowledge questions about violence in their neighborhood, in order to measure their subject-matter awareness and their certainty with respect to their knowledge about violence in their community.

In the analysis, I find that respondents are significantly more likely to update their beliefs about violence in their community (and their support for increased policing) when they receive information from a social media source rather than a government source. Still, the nature of this updating depends on the uncertainty of the respondent. Among those who receive the information from social media, highly certain respondents are significantly more likely to double down on their beliefs, while less certain respondents are significantly more likely to change preexisting beliefs. I observe no such heterogeneous effects among respondents who receive information about violence from the government. Instead, even people with uncertain prior beliefs stick with them when they receive new information from official sources.

These findings suggest that uncertainty only matters when the source of new information is considered trustworthy. To further explore why social media is considered more trustworthy, I conduct a conjoint survey experiment, where I show respondents two pieces of social media information (one related to violence and one to government services) four times each. I then vary the attributes tied to this information by: (1) source, ranging from experts to community members; (2) speed, meaning the time between events and information delivery; (3) post engagement in the form of likes, comments, etc., evaluated here as a signal of

community-driven support for the message; and (4) additional evidence associated with the new information, such as videos or photos. The question asks respondents "which profile would you trust more?" I find that across the board, respondents are more likely to trust information that does not come from the government, is provided in real time, has higher levels of community-level support, and includes additional evidence – attributes tied to social media information. These findings are consistent regardless of the uncertainty levels of respondents and regardless of whether the information is about violence or the provision of government services.

Chapter 5: The Bias of the Crowd

Given that according to my survey experiments, citizens are more likely to change their beliefs when information comes from social media, and that these effects seem to be driven by greater trust in sourcing, among other attributes, the goal of this chapter is to understand what types of information everyday citizens share, and whether it generates a better understanding of the world than other information sources. This is important because the political preferences of citizens largely depend on their beliefs about the real world. In this chapter I investigate how the transition from traditional to social media sources for information influences individuals' beliefs about their reality, with a particular focus on crime. I ask two questions: (1) are social media users more likely than mainstream media to share content related to violence? And if so, (2) does this content accurately reflect official reports about violence? To answer these questions, I draw on nearly 700,000 Facebook posts from 221 local community pages and 1.3 million articles from 21 major media outlets to evaluate patterns in the dissemination of information about violence. I then compare this content to official government crime reports.

In the first part of this chapter, I explore the relationship between community pages and mainstream media content. I argue that social media users will be more likely to disseminate and boost content about local violence on community pages, given that these pages are "closer to home" than major media outlets, which often do not report on isolated incidences

of violence. This expectation is due to the localized nature of the community page, which is specifically designed to provide information about a small neighborhood geography, whereas major media outlets often cover larger geographic areas, if not the entire nation. Consistent with this expectation, I find that although content about violence is not uncommon in the mainstream media, Facebook users are significantly more likely to disseminate and boost information about violence from community pages than from mainstream news outlets. Although traditional media may frequently disseminate sensationalist content, they do so far less than social media users. This is not particularly surprising given the geographic scope of community pages, which are well-suited to sharing information about local violence, due to their hyper-localized nature.

Given that users are more likely to disseminate and boost content related to violence through community pages, I then ask if the content in community pages accurately reflect patterns of violence across the city of Rio de Janeiro. While we might expect an incident of violence to generate corresponding content in community pages, such that more incidences of violence will generate more content related to that violence, I argue that online community pages will diverge from real patterns of local violence for two reasons: (1) poorer communities with higher crime rates may be somewhat desensitized to individual violent events and thus consider them less newsworthy; and (2) middle income communities, which do experience some violence, may be less certain about crime in their neighborhoods than individuals in more or less violent communities, which may lead them to overweight the potential for future violent events, and share each incident with high interest.

When compared to official violence statistics, I find that while better off communities experience significantly fewer violent incidents, posts about violence on Facebook community pages are more numerous in wealthier areas. This pattern suggests a disconnect between the spread of information about violence and actual incidents of it. In evaluating where these patterns diverge, I find that communities in the mid-development levels are significantly more likely to share and react to posts about violence online. Using data from an original survey

of 1,245 respondents, I provide suggestive evidence that respondents from mid-development level communities are more likely than those from low-and upper-development level neighborhoods to be concerned about their safety, defined here as the top and bottom third in terms of social development by neighborhood.

Taken together, the evidence suggests that although people trust social media more because it comes from people "like them," this feature also leads to biased information, colored by the interests of those who create and curate the content. Unlike prior mediums for information, which were also not totally impartial, these biases are inconsistent and not a systematic feature of the editorial slant of the outlet, but rather shaped by the interest, whims, and even fears of millions of distinct users. With respect to violence, these biases may include heightened fear about crime, which then colors the content that is created, thereby raising overall perceptions of violence in a way that may not accurately reflect the real world. I argue that in such contexts, these perceptions may be enough to significantly impact policy decisions and lead to support for more hardline policing approaches that further reinforce the persistent problem.

1.4 Contribution

Violence and conflict affect millions of people around the world. In some cases, the Internet shapes how citizens participate in this reality by allowing them to share their experiences in real-time, whether by organizing protests, facilitating neighborhood watch groups, or reporting on violence. In many countries, social media platforms have served as a vital way for citizens to circumvent established communication channels and shed light on information that might otherwise be ignored due to political pressure or limited journalistic capacity. From contexts as geographically diverse as Brazil, China, the United States, and Egypt, social media has played a vital role in exposing institutionalized violence, police brutality, and a reality of heightened insecurity. In many of these contexts, this type of information

has made previously unknown information known, sparking protest movements or agitation for political reform. It is impossible to know what might have happened without the ability to quickly disseminate this type of content to the crowd. What is clear is that with the global expansion of the Internet – and increased network connectivity through social media – these once neglected narratives are now more difficult to hide.

In some cases, however, communication via social media platforms may heighten racial prejudices or lead individuals to support more militant policing practices out of unfounded fear or anxiety. In recent years, academics and policymakers have viewed social media platforms as a weapon, due to their use during the 2016 election in the United States to disseminate fake news (Allcott and Gentzkow, 2017; Singer and Brooking, 2018). Similarly, during the 2018 elections in Brazil, YouTube's algorithm was optimized for far-right video viewership in support of now-President Jair Bolsonaro. The underlying assumption in both of these cases is that the information provided through these platforms played a significant role in moving public opinion to support certain politicians.

Given the divergent narratives, my research seeks to investigate the explosion of citizen-driven communication through the Internet to evaluate the effect of this new type of information provision on voter perceptions and political attitudes. This is particularly important as the Internet continues to expand around the globe and more citizens begin to rely on social media platforms as their primary source of news. Social media platforms not only provide a space for political manipulation, but also serve as a vital source of information about topics frequently neglected by more traditional forms of media. As a result, it is critically important to evaluate how information transferred through social media platforms transforms political processes and political institutions.

From a theoretical perspective, my dissertation makes an important contribution to understanding the effect of social media on public opinion and voter behavior by focusing on the question of prior beliefs and how the "amateurization" of the information environment affects those beliefs. From a substantive perspective, the proposed research is one of the

first to study how public opinion formation operates in the social media age. This is a new focus of study for political scientists, who have only recently begun to examine the political effects of social media, with a clear focus on the effect of misinformation and polarization. Due to the limited scholarship on the subject, from a theoretical perspective, my research unites disparate, but related literatures from Political Psychology, American Politics, and Comparative Politics.

Past research in Comparative Politics that seeks to assess how changes in the information environment shape voter behavior and beliefs largely neglects the relevant literature in American Politics related to information processing and attitude formation. Few studies that draw on Bayesian updating for theoretical insights about information processing and attitude formation incorporate the strength of prior beliefs into their analyses. I place this variable front and center by arguing that it is a vital factor in Bayesian updating and one that helps to explain the divergent and frequently null findings in the Comparative Politics literature on information and accountability.

2

All the News that's Fit to Tweet

The Internet and associated social media platforms have fundamentally reshaped the way citizens across the world interact with and disseminate information. In turn, this information has altered and transformed a wide variety of political processes. Yet this type of evolution in the communications environment is far from new, with implications dating back to the earliest years in human history. In the first part of this chapter, I briefly sketch the evolution of the information and communications technology (ICT) landscape and highlight the effect of technological changes on political behavior throughout history. I focus on how specific shifts in the information environment (e.g., the volume of the information or the role of intermediaries) may be uniquely suited to altering the political status quo. I then highlight how the internet has simultaneously modified existing processes in and distinctly changed the communications landscape with the introduction of social media platforms. These platforms are defined by their dynamic processes and represent the first ICT to transform information dissemination into a multi-way conversation. I argue that this has unique implications for political behavior, particularly in contexts where traditional information providers are viewed as untrustworthy.

The second part of this chapter delves more deeply into where social media platforms are growing in popularity and highlights unique features about these areas that may have

implications for politics.Yet these are precisely the places where usage of these platforms has stagnated. Instead, it is in countries that are less democratic, younger, and have lower levels of development and press freedom where people are more likely to use social media platforms for news. Looking more closely at the Western Hemisphere, I find that citizens in several South American countries are more likely to use social media for information if they are more interested in politics but trust the government and the media less. By contrast, in the United States (and Canada) citizens are more likely to use social media for information, if they trust the government and media more. This suggests that current political science research focused on the political implications of social media and based in more developed democracies may not be applicable to the political environment in countries where these platforms are growing in popularity, and where this usage is highly correlated with distrust in the government or media. As a result, the conclusions drawn from existing studies may be less relevant outside of a U.S. based context.

For these reasons, it is critical to evaluate the effect of social media on political behavior outside of the United States. In particular, I focus on Brazil, which is far more representative of the average country with respect to development and political and media freedom. Brazil also has a history of high social media usage, which many observers see as responsible for shaping political outcomes in recent years, including elections.

2.1 The Evolution of Information

Since the beginning of human history, individuals have remained focused on the challenge of how to acquire and store more and more information (Gleick, 2011). According to historian Yuval Noah Harari, language first evolved as a way for *sapiens* to share information about the world. However, as the size of groups grew to include more and more people, the potential volume of information generated throughout their network also expanded significantly. In order for humans to grow the size of their communities – a prerequisite for safety and survival

– they had to find a way to keep track of information, in order to be able to recall it when necessary. As a result, they continually had to adapt by identifying new ways to document and store information so that it could easily be accessed at a later date (Harari, 2014).

In early civilizations, this took the form of text and images carved into stone or on parchment or talking drums that in their rhythms stored a message akin to morse code (Gleick, 2011). These now-historic relics and traditions represent some of the first efforts to organize societies and to track an ever-growing volume of information inevitably produced when an expanded number of humans interact amongst each other and form civilizations. For most of human history, the ability to write and read stored information was a craft limited to only a few. But the Industrial Revolution inadvertently opened the floodgates – democratizing the consumption of information and dramatically increasing its production. According to Buringh and Van Zanden (2009), from 501 AD to 1800 AD, the number of books per million inhabitants in Western Europe grew from .1 to 1,000,000. Between 1600 and 1818 alone, the technology fueling this dramatic expansion improved so significantly that publishers went from printing just 240 pages per hour to over 2,400. Never before had it been so easy to document, store, and share information. In reflecting on the effect of movable type, the technology that made mass printing possible, Kuth (2006), noted, "In the centuries that followed, social and legal systems adjusted... and books, newspapers and magazines began to circulate widely." With the printing press, "the age of mass media had arrived."

Since the advent of the printing press in the late fifteenth century, other technological developments have sparked similar exponential growth in the production and consumption of information. The radio, for example, introduced a new means of communication that directly entered into homes on a 24-hour cycle. In 1932, when U.S. President Franklin Delano Roosevelt took office, only 41 percent of Americans had access to the radio. By 1937, more than 90 percent boasted regular access to around the clock radio coverage. Similarly, technological innovations related to the computer dramatically expanded the amount of

information that could be created, processed and shared. In 1951, computer storage could hold approximately 50KB (or five Microsoft World files) on a magnetic drum. Today, we are able to create and store 2.7 zettabytes (or 2,700,000,000,000,000,000 KB) in the "cloud."

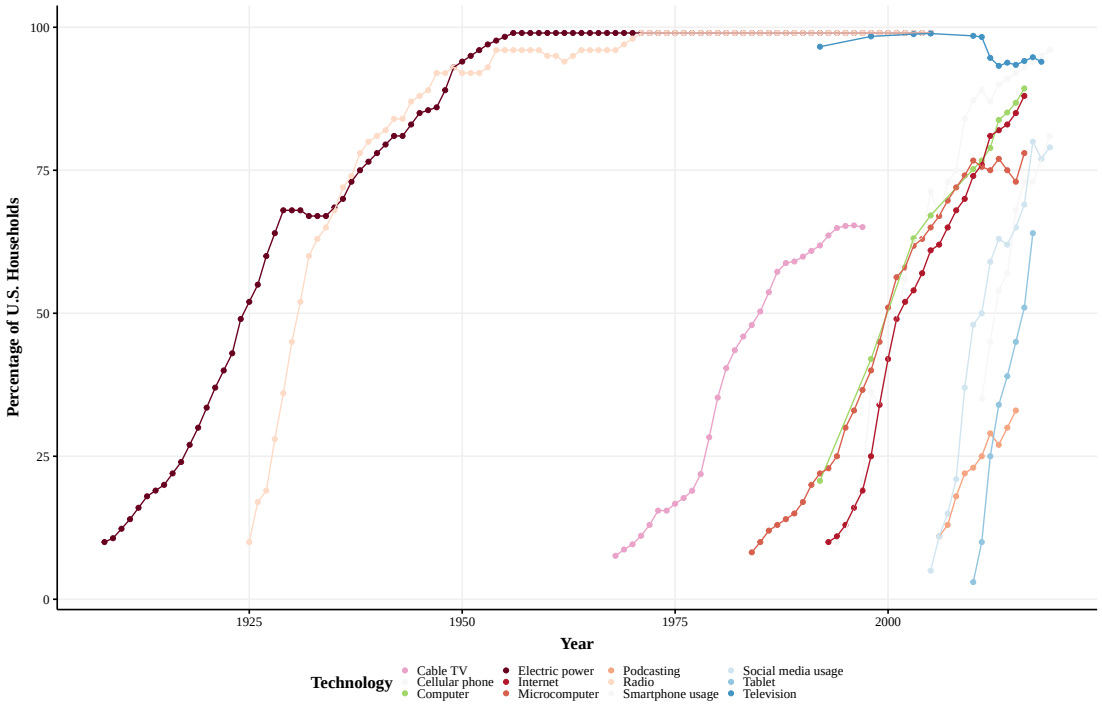

Figure 2.1: Adoption of New Technologies over Time - This figure shows the exponential adoption of twelve communications technologies in the United States over time, as a percentage of total U.S. households. Data on technology adoption comes from Our World In Data.

Throughout history, the rapid expansion of technological capabilities has led to widespread adoption of new technologies by the general population. In **Figure 2.1**, I plot the household adoption rates in the United States of twelve communications technologies over time. While data restrictions prevent me from doing the same for other countries, data from the World Bank on the expansion of Internet access across the world, plotted in **Figure 2.2**, also mirrors these patterns. Equally important, Internet adoption across the world is likely to further accelerate as more options become available to less developed areas through satellite and other services.

Disruptive changes in the information environment throughout history have triggered

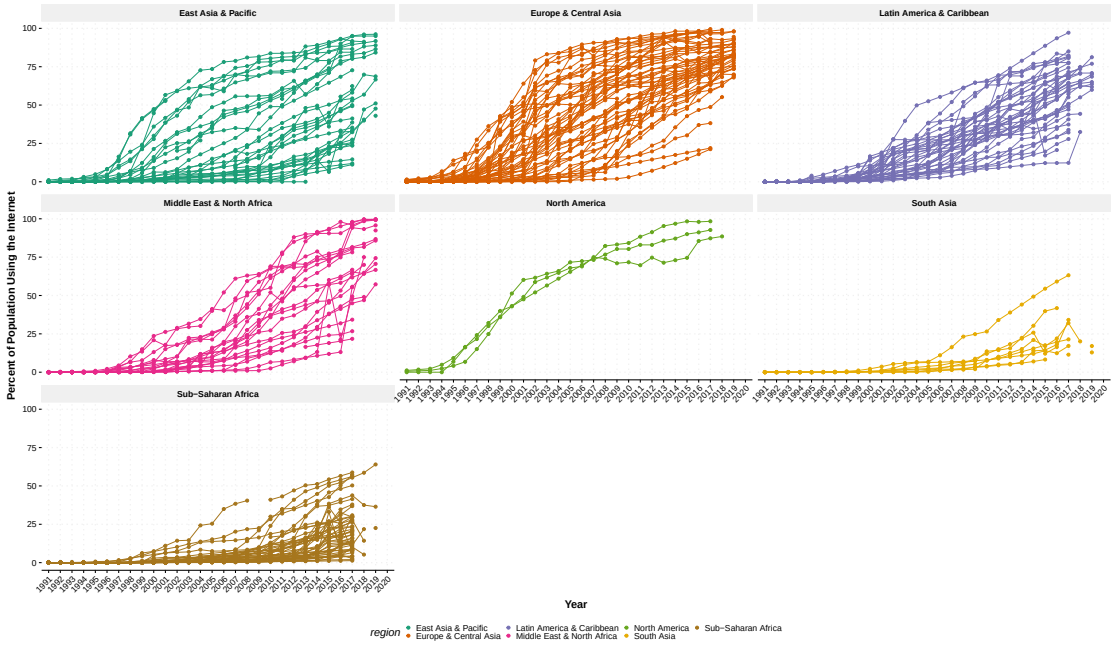

Figure 2.2: Percentage of the Global Population Using Internet - This figure shows the exponential adoption of the internet in 266 countries around the world over time, by region. Data on the percentage of the population using the Internet comes from the World Bank.

equally disruptive transformations in political systems. One obvious manifestation of this dynamic is the way in which the printing press sparked an entire religious movement, catalyzed by Martin Luther's publication and mass distribution of the *Ninety-five Theses* in 1517. This document, along with others that followed, led to a split in the Catholic Church and sparked the Protestant Reformation, which fundamentally altered political and religious dynamics across Europe, with implications throughout the world. More broadly though, the printing press ushered in an era of mass communication by lifting limitations on how information circulated or where new ideas might take root and subsequently spread. Over the centuries, education and literacy ceased to be privileges of the elite. New ideas – including new scientific discoveries – could be widely disseminated to the masses. This led to further technological innovation, longer life spans, and a better informed public playing a more critical role in shaping political and economic outcomes.

In the United States, the "Golden Age of Radio" likewise ushered in a new era of political engagement and allowed presidents to communicate directly with the public for the first time

ever. The rise of radio programming also marked the first time that information was provided to consumers on a 24-hour-cycle. During his fireside chats, President Roosevelt expertly used the radio to calm an anxious public grappling with the Depression and the chaos of World War II. Through broadcasted addresses, Roosevelt spoke directly to voters during difficult political years and encouraged them to believe that they had a friend in the White House (Levine and Levine, 2010). Importantly, these conversations cut out the intermediaries – such as journalists or business executives – and allowed the president to speak directly to the public in a manner similar to how politicians use platforms like Twitter today.

Across the Atlantic Ocean, the radio was also used for more sinister ends, becoming a vital tool in the dissemination of Nazi propaganda. Joseph Goebbels, who served in the role of Reich Minister of Propaganda for twelve years, was the mastermind behind the Nazis' propaganda strategy and pioneered the mass production of the "Volksempfänger," or the "People's Receiver," which had the tagline "All of Germany hears the Führer with the People's Receiver." According to Adena et al. (2015) the widespread availability of relatively cheap radio access in Nazi Germany played a vital role in shaping citizens' support for their political leaders, including Adolf Hitler. By 1941, 65 percent of the population owned one, which brought propaganda directly into the homes of everyday Germans (Rentschler, 2003). Similarly, a study of the spread of 3G cell phone networks, Guriev, Melnikov and Zhuravskaya (2019) found that the that the expansion of Internet access through cell phones between 2008–2017 increased voters' awareness of corruption and reduced their support for incumbent governments in countries without Internet censorship.

What is important to note here is that each subsequent technological change dramatically altered the way citizens acquired, processed, and stored information about their surrounding environments. Citizens would then use this newly available data to make decisions about their political attitudes and actions. Up until recently, the masses benefitted by receiving this information, and deciding what to do with it moving forward. This one-way street, while transformative in its own right, did not give the masses the ability to talk back to elites or

to communicate with eachother in real-time about the information they received. In fact in many ways, it handicapped audience participation, relegating the public to passive consumers who could only respond through a letter to the editor. However, the latest, seemingly seismic, technological shift in the information environment has transformed this one-way information flow into a multi-directional conversation. Not surprisingly, the expansion of the Internet and associated social media platforms over the past few decades has once again transformed the information environment. Their participatory nature however is fundamentally different from past communications innovations, and the role of this dynamic in shaping political outcomes has rapidly emerged as the latest frontier in academic research.

2.2 What is Social Media?

Although technological change throughout history – ranging from the printing press to the television – rapidly expanded the speed and volume with which humans could receive, process and store information, much of what was shared through these mediums was a one way street, where citizens served as recipients, and elites – in the form of publishers, news anchors, editors, etc. – served as the main content creators. The process by which information was created and disseminated took place within a professional framework, and citizens had little say in what was shared beyond decisions about whether or not to tune in (Shirky, 2008). As a result, this information was filtered by what a select (often highly educated) few deemed both worthy of sharing and capable of generating profits. The expansion of social media to the masses, according to Jensen (2015), has fundamentally altered this "one-to-many" dynamic, replacing it with a "many-to-many" information environment where everyday citizens are active contributors in shaping the conversation. The rapid growth of these platforms is plotted in **Figure 2.3**. This exponential growth mirrors that of other communications technology innovations highlighted in previous sections.

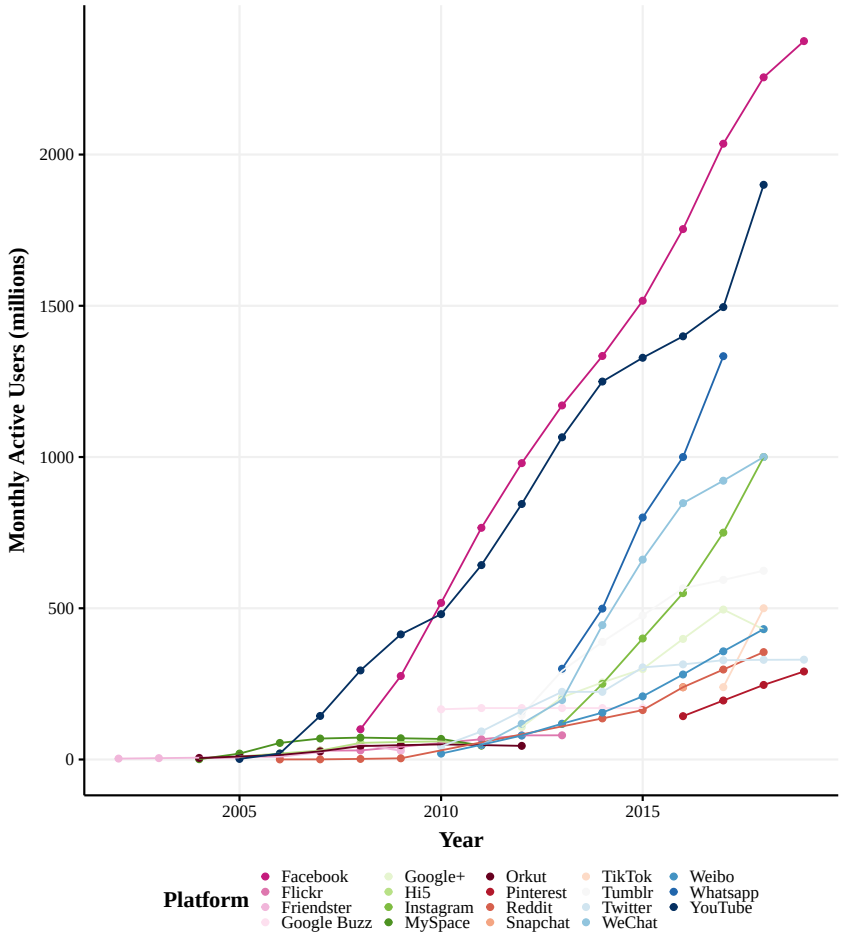

Figure 2.3: Growth of Social Media Platforms Over Time (by Platform) - The figure shows the global growth of nineteen social media platforms over time as measured by Monthly Active Users. Data for this figure comes from *Our World In Data.*

DiNucci (1999) is commonly credited with the first major discussion of how the evolution of Internet capabilities could transform the flow of information from a one-way monologue into a multi-directional conversation. Writing in 1999, she predicted that:

> "The Web we know now, which loads into a browser window in essentially static screenfuls, is only an embryo of the Web to come. The first glimmerings of Web 2.0 are beginning to appear, and we are just starting to see how that embryo might develop. The Web will be understood not as screenfuls of text and graphics but as a transport mechanism, the ether through which interactivity happens. It will...appear on your computer screen,...on your TV set...your car dashboard...your cell phone...hand-held game machines...maybe even your microwave oven."

This interactivity came to be the defining feature of "Web 2.0," or the "social web", with social media as a central player in this reimagined Internet environment. As Kaplan and Haenlein (2010) describe it, social media is: "a group of Internet-based applications that build on the ideological and technological foundations of Web 2.0, and that allow the creation and exchange of user-generated content." Bechmann and Lomborg (2013) further elaborate on the distinctions between social media and other traditional forms of information provision, namely radio or television. They argue that social media has three defining characteristics that distinguish it from prior information generating platforms: (1) communication is de-institutionalized; (2) the user is regarded as a producer; and (3) the platform is interactive and networked.

Additional research has pointed to other attributes of social media that may distinguish it from prior information mediums, including the speed with which information spreads (Nahon et al., 2011), the potential for it to get lost in the noise and become "digital dust" (Margetts et al., 2015), and the continued shift toward sloganeering and performative discourse (Gillespie, 2010). I gather some of these features into **Table 2.1**. In this table, I highlight the various facets of information and communications technology, how these features operate in

traditional media, how they might change with respect to social media and the implications this shift may have for politics.

Table 2.1: Political Implications of the Difference between Social and Traditional Media

Feature	Traditional Media	Social Media	Political Implications?
Audience	Targeted to user (e.g., media market)	Algorithmically-determined, but in theory, anyone	Algorithm exposes users to only to like-minded people (sorting); Much broader exposure to ideas and information
Speed	Limited by capacity of outlet	Real-time or perhaps, near real-time (by capacity of users)	Quicker reactions to events; Viral exposure and potential mischaracterization/distortion
Volume	Limited by capacity of outlet	Unlimited (or perhaps, limited by capacity of users)	Previously unseen information can now be captured/ documented; Oversaturation
Info Flow	One way flow	Multi-directional conversation	Improved accountability/dialogue between elites and population; Elite distortion/citizen correction; Easier coordination
Creators	Executives, editors, fact-checkers, journalists (elites)	Anyone with internet access (non-elites)	Rise of amateurism; Greater potential for mischaracterization/distortion
Content	High level of detail	Low level of detail	Preference toward short soundbites

Table 2.1 provides an overview of the features that distinguish social media from traditional media (television, radio, newspapers) and highlights the potential political implications of these differences.

While this table is by no means exhaustive, it captures some of the ways in which heavy reliance on social media as a primary source of information might generate distinct political outcomes. We should expect citizens who are more reliant on social media for information about the external environment to be more susceptible to this type of political effect. In other words, where are citizens more likely to use social media beyond its social capabilities as a (or the) primary source of information? Drawing on cross-national data, I explore this question in greater detail in subsequent sections.

2.3 Social Media As News

Much of what we know about how social media affects political attitudes is shaped by a small body of experimental and survey research geographically anchored to the United States and other developed democracies (e.g., Bail et al., 2018; Allcott et al., 2019; Lee et al., 2014; Allcott and Gentzkow, 2017). Although each of these studies has contributed significantly to our understanding of the role of the Internet in shaping political outcomes, their geographic limitations are problematic for a variety of reasons. As Jungherr, Rivero and Gayo-Avello (2020) note:

> "The United States is also very different from any other democracy. For one, elections are a multi-million-dollar industry. This has led to the emergence of a rich ecosystem of consultancies willing and able to invest in the development of digital tools and services supporting candidates on any level of politics. Over the last twenty years political discourse in the country has degenerated into a veritable blood sport, with the two sides of the aisle pummeling each other mercilessly. This has led to a weakening of political institutions, rendering them highly vulnerable to challenges by digital media" (pg., 7).

Similar social science research outside of the U.S. context largely focuses on protest mobilization during the Arab Spring (e.g., Tufekci and Wilson, 2012; Wolfsfeld, Segev and Sheafer, 2013; Breuer, Landman and Farquhar, 2015) or on autocrats' use of these platforms for political purposes (e.g., Sobolev, 2018; King, Pan and Roberts, 2013). Only a few recent studies explore the broader political effects of social media usage in a non-U.S. based context (e.g., Asimovic et al., 2021). Although this research has generated valuable insights, the exponential growth in day-to-day usage of these platforms has occurred primarily outside these contexts or unique events. In a *Pew Research Center* study, Poushter, Bishop and Chwe (2018) found that although social network usage has plateaued in a number of developed economies, it has continued to expand rapidly across many less developed countries.

Specific to Facebook, the social media platform on which I focus my analysis in **Chapter 5**, the growth in usage of this platform has been constant and rapid, despite the emergence of other competing social networking sites (**Figure 2.3**). However, similar to the findings in the Pew Report, this growth has largely stagnated in Europe and the United States, and is now primarily concentrated in other parts of the world.[1] The growth of Facebook's monthly active users, defined as users who sign onto Facebook at least once per month, is plotted in **Figure 2.4** using the geographic designations internal to Facebook to describe the location of its users in the company's annual reports to shareholders. Growth in active Facebook users has hardly risen over the years in the geographic regions where the effect of these platforms is studied the most.

Importantly, it is the areas that are seeing the most growth in social media usage where these platforms are used beyond their "socializing" objectives and as a primary means for gathering information and news. In these countries, for example, there is a significant negative correlation between HDI and the use of social media for news. Specifically, a one unit increase in HDI is associated with a .75% decline in social media usage for news (**Figure 2.5, top right**). While the sample is restricted to only the forty countries surveyed by the Reuters Institute, these results suggest that social media usage for news purposes is more likely in less-developed countries, where the media is potentially less established. Data on the percentage of the population that gets their news from social media is drawn from the Reuters Institute's 2019 Digital News Report.

We also see this relationship play out in other important areas. As shown in **Figure 2.5**, the percentage of the population that obtains news from social media is negatively associated with the Economist Intelligence Unit's Democracy Index (top, left figure), which measures the country-level state of democracy based on five factors, including: "electoral process and pluralism, the functioning of government, political participation, democratic political culture and civil liberties." In this sense, in countries that are less democratic, citizens are

[1]Unfortunately, Facebook does not distinguish between regions in its classification of "Rest of World," however, we can assume this includes both Latin America, Africa, and the Middle East.

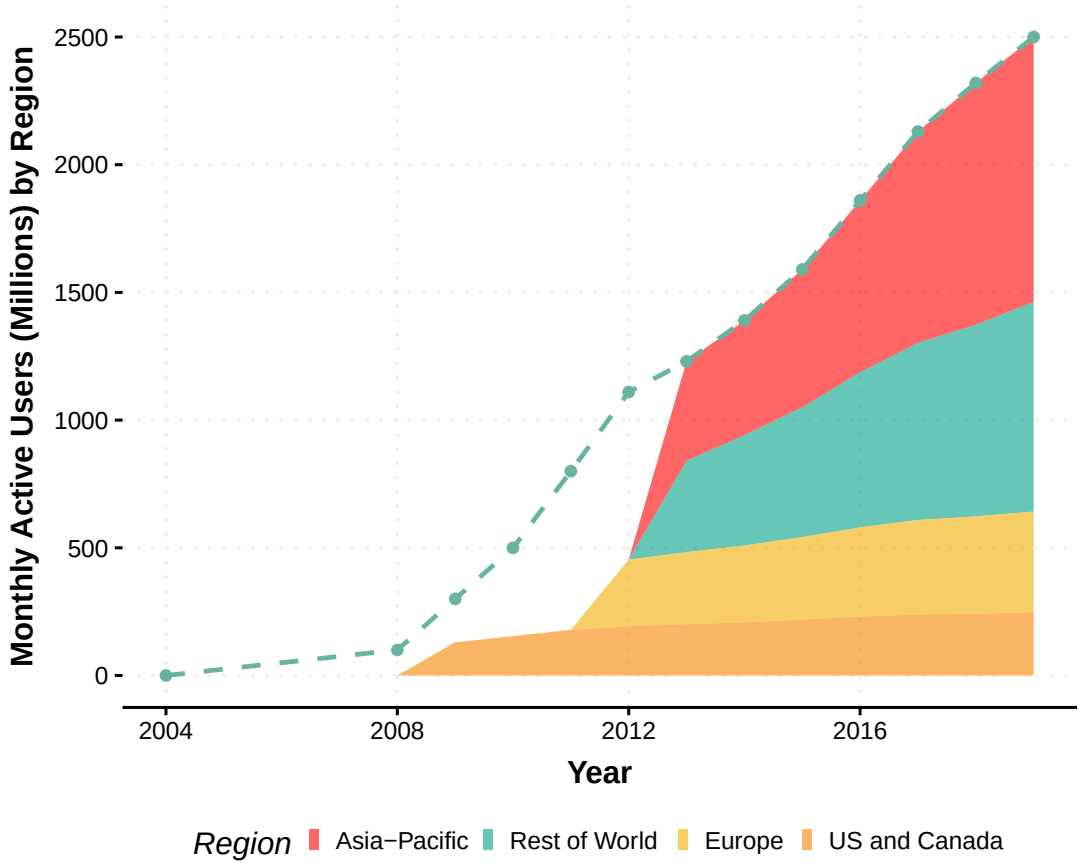

Figure 2.4: Facebook Growth Over Time - The figure shows the growth in Facebook's Monthly Active Users, at both the aggregate and by geographic region. Data comes from Facebook's annual reports to investors.

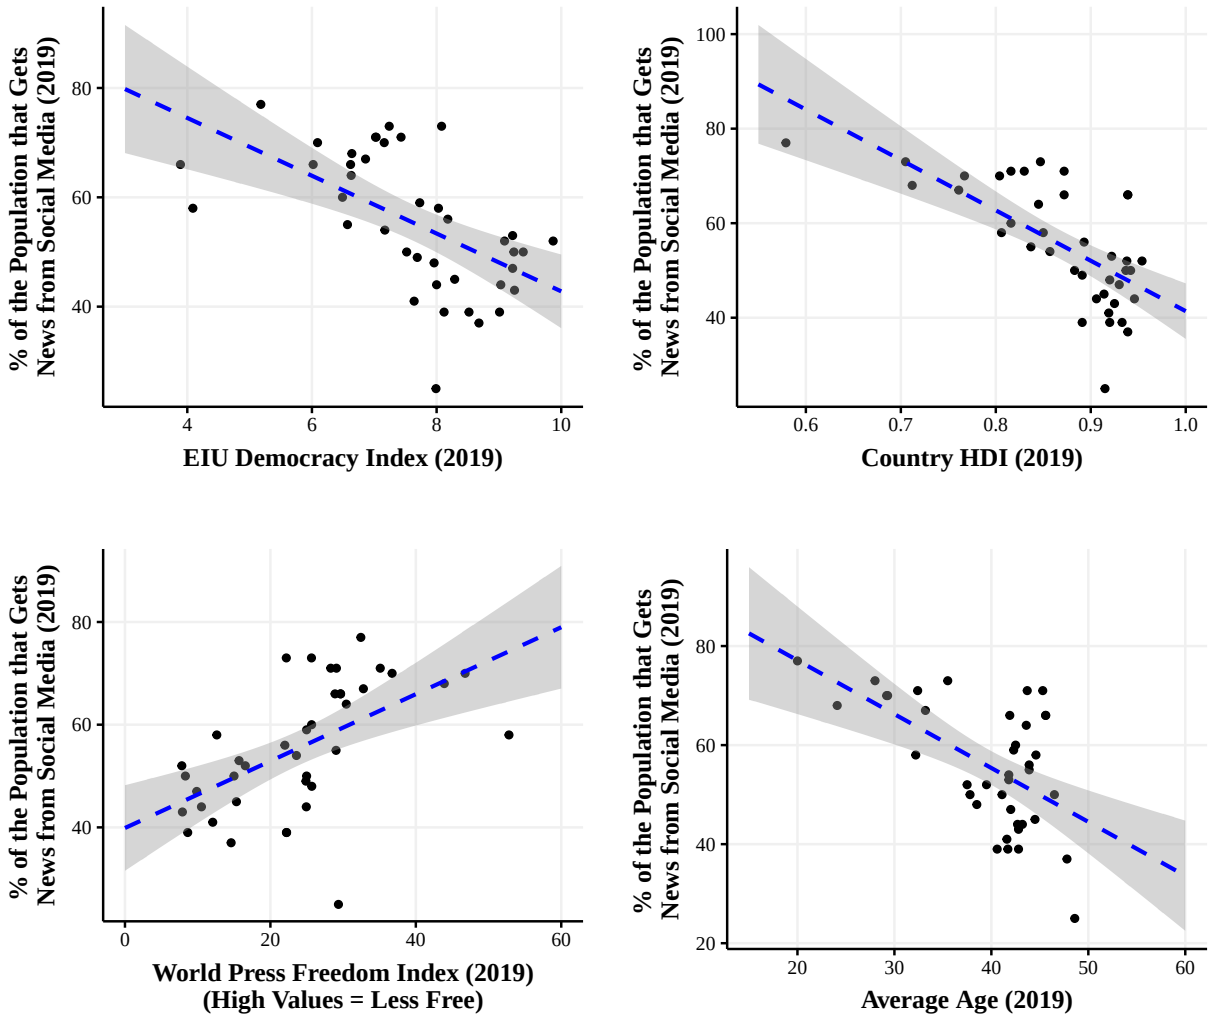

Figure 2.5: Relationship between the Percent of Citizens that Use Social Media for News and Country-Level Attributes - This figure shows the bivariate relationship between the percentage of respondents who use social media for news in 2019 and (1) levels of democracy; (2) Human Development Index; (3) press freedom; and (4) average age at the country-level. Data on the percentage of respondents that report they get new from social media comes from the Reuters Institute Digital News Report.

more likely to use social media for news. Additionally, countries with worse scores on press freedom (bottom, left figure), as measured by the Reporters Without Borders' World Press Freedom Index and with younger populations (bottom, right figure) are also more likely to acquire their news from social media, as explored in **Figure 2.5**. In this figure, a higher press freedom index score is associated with less press freedom.

These results, while preliminary in nature and limited to a small sample for which data are available, demonstrate that current research on the political effects of social media may be missing important dynamics in countries where these platforms are exponentially growing in popularity. In particular, theory-building and empirical analyses centered around social media usage in a small number of highly-developed countries with high press freedom or democracy scores may generate contextually-specific findings that are not easily portable to the distinct contexts where these platforms are actively growing their user base. These country-level differences may also influence other individual-level variables that factor into why people seek out political information on social media. This further complicates generalized theories of political change due to social media exposure, which are largely evaluated using data from U.S. respondents.

Drawing on the Latin American Public Opinion Project's (LAPOP) survey data from 2018-2019, I evaluate in **Table 2.2** the association between using social media for politics and (1) interest in politics, (2) trust in the media, and (3) trust in government in the U.S., as compared to other countries in the Americas. Survey responses come from ten different countries, seven in South America and three in North America, for a sample size of 15,594 survey respondents. Of this total, 8,287 (53%) do not get political information from social media. The remaining 7,307 (47%) do. In this sample, social media includes Twitter, Facebook or WhatsApp. Facebook is the most common platform for political information, with over 84% of respondents who report getting political news from social media doing so from Facebook. The second most popular site is WhatsApp (41%), followed by Twitter (17%). Twitter is also the least popular social media platform, both globally and in the

survey sample.[1]

I use the following regression model with an interaction coefficient, detailed in **Equation 2.1**, to evaluate the relationship between social media usage for politics (independent variable) and government trust, media trust, and interest in politics (dependent variables). I interact the independent variable with a binary dummy variable, which codes respondents with a `1` if they are part of the United States survey and a `0` otherwise. The equation specifications are:

$$\begin{aligned} \text{DV}_i = &\beta_0 + \beta_1[\text{SM Usage}_i] + \beta_2[\text{US}_i] + \beta_3[\text{SM Usage}_i * \text{US}_i] \\ &+ \beta_4[\text{Sex}_i] + \beta_5[\text{Age}_i] + \beta_6[\text{Region}_i] + \beta_7[\text{Income}_i] + \epsilon_i \end{aligned} \tag{2.1}$$

In theory, if there is no fundamental difference between respondents in the U.S. context and respondents outside the U.S. context, we should see similar patterns across interaction coefficients. This would signal that motivations for seeking out information about politics online, in particular, are similar across the board. It would then follow that resulting theories of political behavior, largely evaluated in the U.S., should hold across different socioeconomic and political contexts. Yet, this is not what I find in the survey data.

Models (1), (3) and (5) plot the simple bivariate relationship between usage of social media for politics and interest in politics, trust in media and trust in government, respectively. Survey respondents who use social media for politics are significantly more likely to be interested in politics, less likely to trust the media, and less likely to trust their government. However, this relationship masks a clear divergence based on the geographic location of respondents. Specifically, the interaction effects from these correlations, detailed in **Table 2.2**, demonstrate that survey respondents in non-U.S. contexts who use social media

[1]Survey respondent breakdown by country includes: Brazil (1,498 respondents), Mexico (1,580 respondents), Colombia (1,663 respondents), Argentina (1,528 respondents), Bolivia (1,682 respondents), Ecuador (1,533 respondents), Peru (1,521 respondents), Uruguay (1,581 respondents), United States (1,500 respondents), and Canada (1,508 respondents). Interviews were conducted face-to-face with survey respondents sampled using a probabilistic design stratified geographically to ensure representation across each country, in both urban and rural settings. More detail on the technical information for the 2018/19 AmericasBarometer wave is available here: Technical Information, AmericasBarometer, 2018/19.

Table 2.2: Association between Using Social Media for Political Info, Interest in Politics & Trust

	Dependent variable:					
	Interest in Politics		Media Trust		Government Trust	
	(1)	(2)	(3)	(4)	(5)	(6)
$\mathbb{1}$(SM Use for Politics)	0.458***	0.460***	−0.221***	−0.284***	−0.050	−0.085**
	(0.019)	(0.021)	(0.034)	(0.037)	(0.035)	(0.037)
$\mathbb{1}$(US)		0.637***		−1.230***		0.672***
		(0.076)		(0.135)		(0.136)
$\mathbb{1}$(SM Use for Politics)×$\mathbb{1}$(US)		−0.004		0.500***		0.281***
		(0.056)		(0.102)		(0.102)
$\overline{y_i}$	1.647	1.837	4.873	4.625	3.727	4.012
Observations	12,080	12,080	12,738	12,738	12,526	12,526
R^2	0.154	0.160	0.057	0.063	0.007	0.012
Adjusted R^2	0.153	0.159	0.057	0.063	0.007	0.011

Table 2.2 shows the association between using social media for politics and Interests in politics (models 1 and 2); trust in the media (models 3 and 4); and trust in government (models 5 and 6). All models control for region, sex, age, and income (coefficients not shown). Model (2), model (4) and model (6) interact the binary variable social media for politics with a binary variable for whether the respondent is in the U.S. for each respective dependent variable. Model (1), model (3) and model (5) exclude this interaction. In addition to the U.S., respondents are located in Brazil, Mexico, Bolivia, Colombia, Ecuador, Uruguay, Argentina, and Peru. Significance: *p<0.1; **p<0.05; ***p<0.01.

for news are significantly more likely to be interested in politics (model 2), less likely to trust the media (model 4), and less likely to trust their government (model 6). By contrast, survey respondents who use social media for news in the U.S. are significantly *more* likely to trust the media (model 4) and more likely to trust the government (model 6). They are also slightly less likely to be interested in politics, but these results are not statistically significant. In all models, I control for region, gender, age and income.[1]

Due to data limitations, I am unable to incorporate countries outside of North and Latin

[1]I use the following question to measure "Interest in Politics": "How interested are you in politics?" (*Very (4), Somewhat (3), A Little (2), Not at All (1)).* I use the following question to measure "Trust in the Media": "How much do you trust the mass media?" (*Response vary 1-7, with one being not at all and seven being completely.*) I use the following question to measure "Trust in the Government:" "How much do you trust the legislature?" (*Response vary 1-7, with one being not at all and seven being completely.*) To measure "Social Media Use for Politics," I combine responses to three questions that ask whether respondents use social media for politics for Twitter, Facebook or WhatsApp. If respondents respond that they use one of these platforms "Daily" or "Several times a week", they are classified as "using social media for politics", otherwise they are classified as not using social media for politics.

America into this analysis or explore additional variation by year; however, these results point to a potential challenge to the reliance on existing studies to form general theories about how social media shapes political outcomes world-wide. What is clear based on these results is that the motivations driving people to these platforms for news and politics may be different in different contexts, and in particular, in contexts where the population is younger, development is not as high, democracy is less "established," or the press is considered less free. Given these associations, it is important to understand the patterns of usage and the political ramifications of these platforms in contexts with less-established free press environments and in countries that are less developed, but where the user base for social media platforms is growing rapidly.

2.4 Social Media Usage in Brazil

In an effort to expand the geographic scope of current social science literature on social media's political effects, the geographic focus of my dissertation is Brazil, where out of a total population of 211 million, nearly 150 million people use the Internet, and as demonstrated in **Figure 2.6**, a majority of Internet users get their news primarily from social media. Importantly, Brazil consistently ranks in the top five countries among active Facebook users, with over 120 million users, behind just Indonesia, India and the United States – all of which have larger populations. With respect to Facebook penetration, Brazil is tied for first with the United States in terms of Facebook users as a percentage of the total population (57%). Focusing on both Press Freedom and HDI, Brazil falls squarely in the middle, with a press freedom score that is exactly at the mean of 34 and a country-level HDI that is slightly above the mean of 0.7149, at .761; however, among countries that are also at or close to the mean with respect to measures of press freedom, this deviation from the average country-level HDI is the smallest.

Finally, another reason to study the political effects of exposure to social media as news

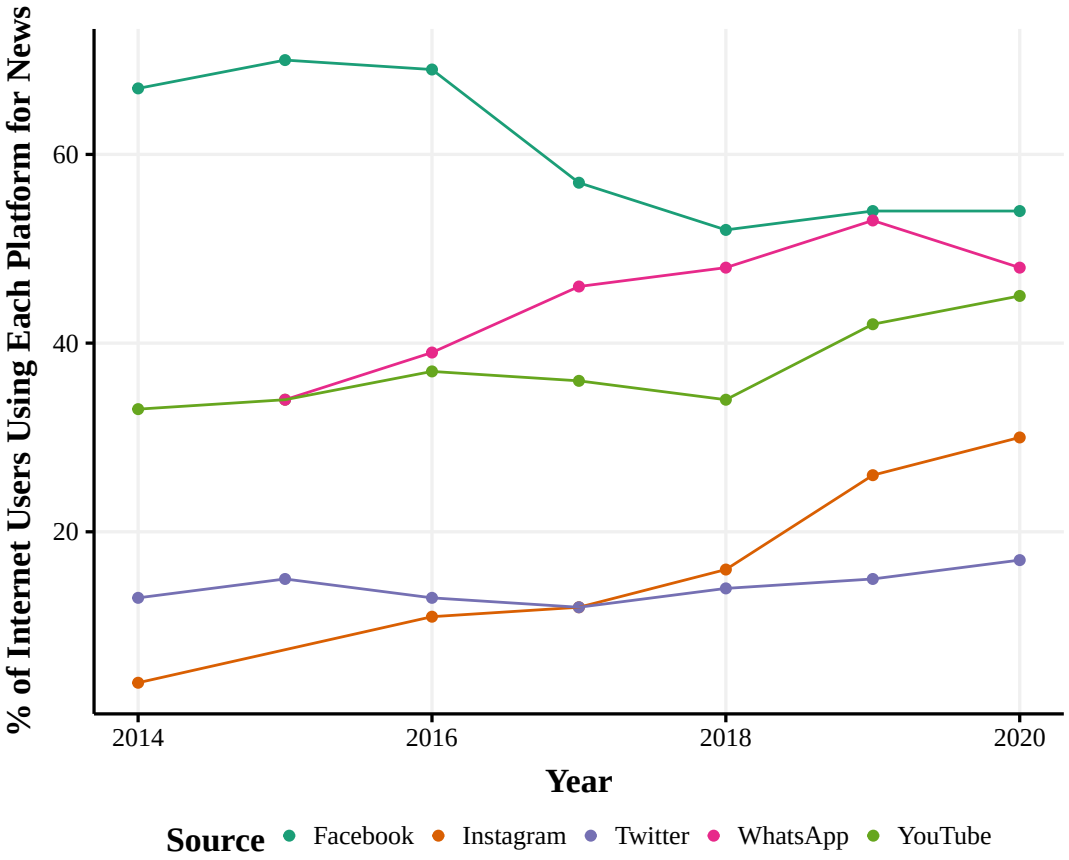

Figure 2.6: Percentage of Brazilian internet-users who get their news from social media. - Annual number of survey respondents that report they get new from social media (by platform) in Brazil, according to the Reuters Institute Digital News Report.

in Brazil is that the country's 2018 election was marred by accusations of social media biases, which may have helped to sway the election toward the right-wing, populist candidate Jair Bolsonaro. From YouTube to WhatsApp to Facebook, story after story reports on the influence of these platforms in pushing Brazil's population toward far right politics and disseminating false or misleading content. While clearly this is a problem, as in the United States, it is hard to know exactly who is most susceptible to this type of (mis)information. Given the speculative association between social media information and political outcomes in Brazil, the "average" nature of the country with respect to both development and media freedom, and the volume of people that use these platforms (and also use them for news), Brazil is an ideal country to evaluate the relationship between social media information and public opinion formation.

In Brazil, social media is also a major source of local-level information. Citizens use these platforms to communicate about political issues at the community level, including city services, violence, and community events, and they share stories of community resilience (See: CESeC, 2019). In **Chapter 5**, I analyze the content of these community-level pages in more detail. However, in an original representative survey of 1,246 Rio de Janeiro residents (which I analyze in depth in **Chapter 4**), I find that more than half (56%) of all respondents in this major metropolitan area of nearly seven million participate in community-focused pages or groups (**Figure 2.7**).

Furthermore, even if they are not formally members of an online neighborhood community, they are much more likely to cite social media as their main source of local-level information. This is illustrated in **Figure 2.8**, which plots the frequency with which respondents turn to social media (WhatsApp and Facebook) vs. traditional media (television and newspapers) to get information about their community. For this question, respondents were required to rank potential sources of community information based on which they use the most and least frequently. Disaggregated, respondents overwhelmingly rank Facebook and WhatsApp in their top two, followed by television and newspapers.

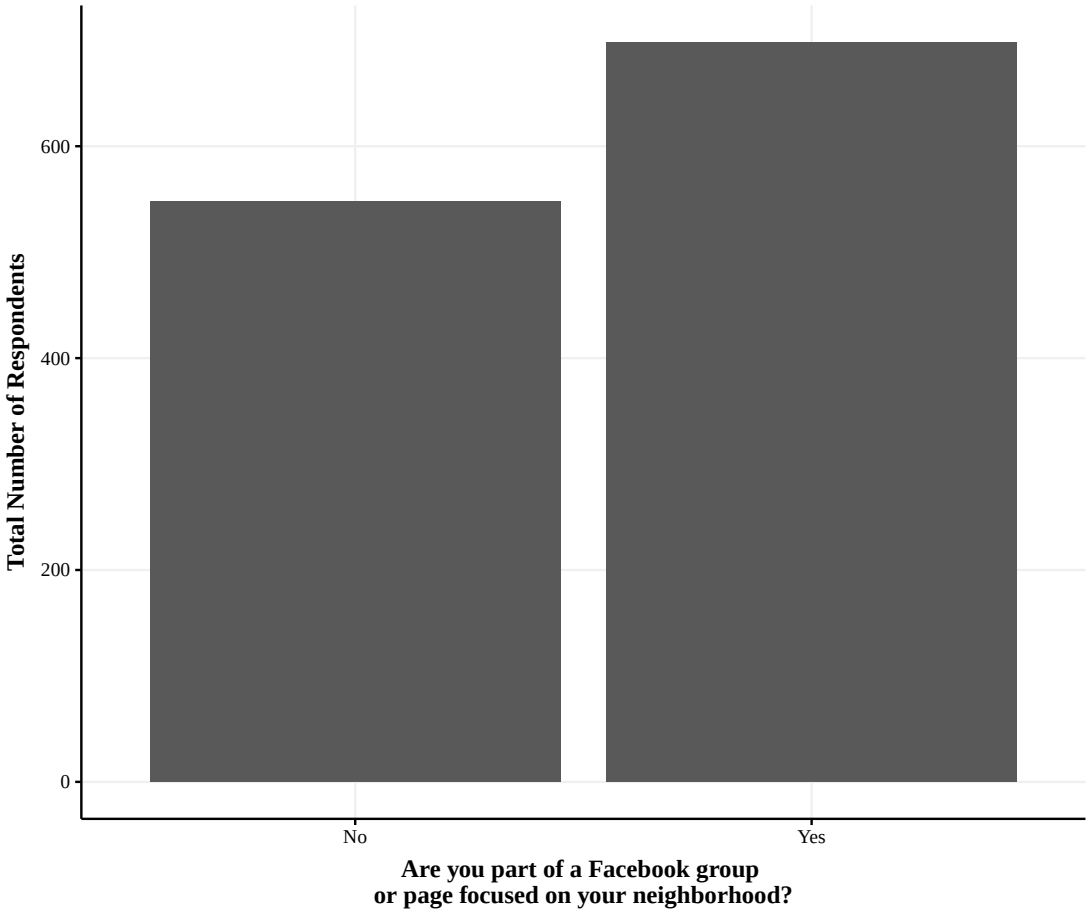

Figure 2.7: Community Facebook Group and Page Membership - Number of survey respondents who report being members of a community Facebook page.

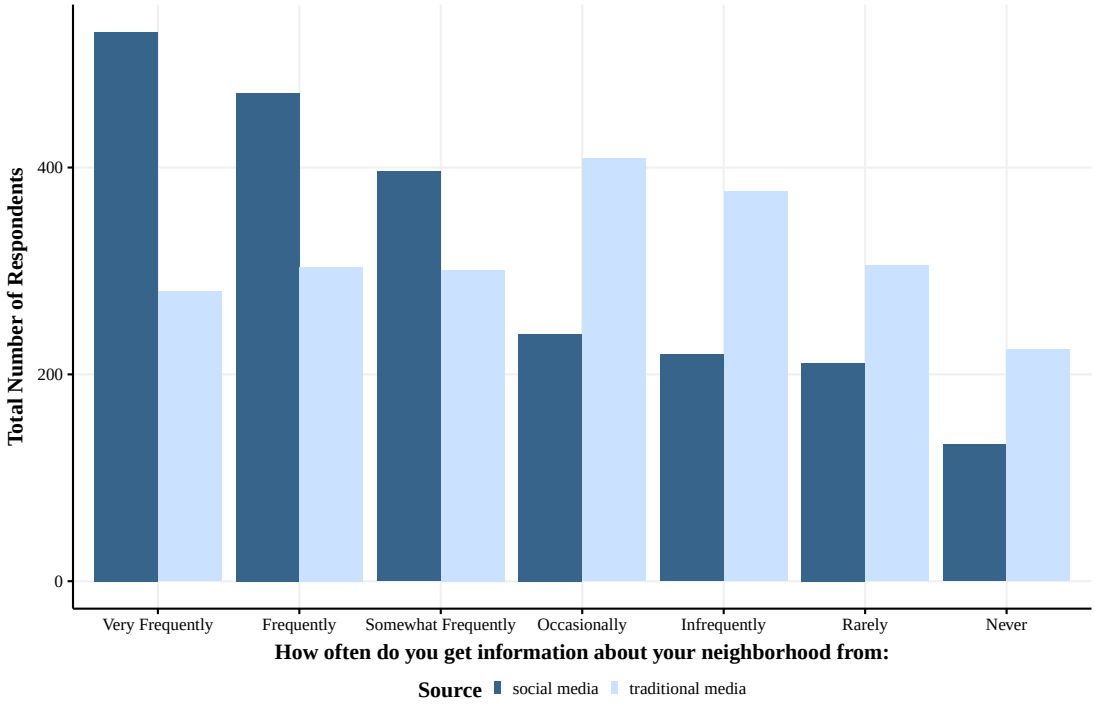

Figure 2.8: Frequency of Neighborhood Information by Source - Survey respondent rankings of source of information about their community. Social media includes WhatsApp and Facebook. Traditional media includes television and newspapers.

What is evident from these survey results is that in Brazil (and particularly in Rio de Janeiro), citizens have begun to expand the uses of social media beyond its original intent of serving as platforms for connecting and sharing information about peoples' lives. In contexts such as Brazil, citizens are turning to these platforms as a way to gather vital information about their day-to-day existence, supplanting traditional information mediums like the local news, newspapers or the radio. The reasons for this shift may have to do with the fact that citizens in this context are less trusting of official channels of information (those that follow a one-to-many logic) and more willing to believe information that comes from non-elites, or people just like them (a many-to-many logic). I explore this phenomenon further in **Chapter 4**. I also explicitly incorporate this expectation as a central facet of my overarching theory, detailed in **Chapter 3**.

2.5 Conclusion

In this chapter, I have provided a brief overview of the history of the technological landscape, highlighting how new technologies have throughout history transformed political processes as a result of shifts in the information environment resulting from these innovations. For example, with the printing press the speed with which information could be published was dramatically altered, allowing for a more rapid dissemination of information. The radio marked the first time that politicians could speak to the masses on a 24-hour-cycle, diminishing the value of intermediaries like media executives or journalists. With the growth of social media platforms, this direct line to the masses has now shifted to a dialogue, with the masses able to communicate back in unprecedented ways.

These changes, unsurprisingly, have catalyzed equally dramatic political shifts. With respect to social media, however, social science research has to this point paid only limited attention to how these dynamics play out in non-U.S. contexts. Importantly, however, it is these contexts where social media usage is rapidly growing and where the transition from

one-to-many vs. many-to-many content may be critical due to the less established nature of democratic governments or free press. In this chapter, I present preliminary evidence that in these contexts, citizens appear to rely more on social media platforms for news about their surrounding environments. I highlight several trends that distinguish the types of environments where social media usage is growing from the types of environments where the majority of academic research has occurred to date.

In particular, citizens are more likely to get their news from social media in countries that are less democratic, less developed, have lower press freedom, and are younger. By contrast, research on the subject has occurred in countries that are more democratic, more developed, have higher press freedom, and are older. In more developed countries such as the U.S., citizens who use social media for news are more likely to trust the media; whereas in non-U.S. contexts, citizens who use social media for politics are more interested in politics, and less likely to trust the media or the government.

These context-based and motivational distinctions matter when evaluating the effects of social media information, and may play a significant role in shaping a wide range of political outcomes. In the context of this dissertation, I focus on one country – Brazil – and one type of political outcome – support for policing. In addition, I develop a broader theory of change which I expect to be applicable in contexts where political and media environments are less well-established, due to, for example, longstanding conflicts or a history of authoritarianism, as is the case in Brazil.

3

A Theory of Information Processing in the Digital Age

> "The people formerly known as the audience wish to inform media people of our existence, and of a shift in power that goes with the platform shift you've all heard about."
>
> – *Jay Rosen, "The people formerly known as the audience", 2006*

Political science theories focused on public opinion formation have long argued that when exposed to new information, people will update their beliefs, but that the effect of new information may depend on how firmly they hold their initial beliefs. Extending this basic theory, researchers have evaluated the ways that people filter new information to comport with their general worldview or partisan preferences, a phenomenon known as motivated reasoning, and the role of source expertise in shaping the updating, particularly the role of experts. These theories, and the empirical evidence that they produce, were explored prior to the advent of social media. This development in the communications environment has transformed the flow of information from a monologue—where elites speak to the average citizen and citizens either move their opinions closer to those expressed by elites or resist elite persuasion—to

a multi-way conversation—where "the people formerly known as the audience" can easily speak back and amongst themselves to enhance the credibility of information from community sources (Rosen, 2008). This dissertation evaluates whether or not differences between social media and more elite-driven forms of political communication from sources such as the government differentially affect public opinion formation and change, and thus support for different types of political outcomes.

The theory I put forth in this chapter follows from the findings discussed in **Chapter 2** and will be tested in **Chapters 4** and **5**. I argue that two main conditions determine when new information leads to updating: (1) the certainty of prior beliefs; and (2) trust in the sourcing of the information. While the former point is universal across contexts, I argue that the latter is context-specific, varying based on the political environment of those who receive the information.

In brief, I argue that more information, if it challenges prior beliefs and is salient to voters, should in theory lead to a reassessment of prior opinions and catalyze a shift in policy preferences. However, the relationship between new information and prior attitudes is complicated by a second variable – the certainty of an individual's prior beliefs. If someone is highly certain about their prior beliefs, evidence that disconfirms that belief will produce little meaningful updating. If someone is less certain about their prior beliefs, evidence that disconfirms that belief may lead to a change in attitudes. This is the foundation of Bayesian updating.

Weak priors, however, will not automatically induce updating: the nature of the information also matters, a factor that largely goes unaddressed in traditional models of information processing. I argue that while the shift into the digital space altered the information environment in a variety of different ways, many of the changes were a continuation of prior changes catalyzed by previous shifts in the information and communications technologies landscape. What is unique about social media is that it has transformed the information environment from a one-way monologue to a multi-way conversation, where elites can speak directly to

citizens, citizens can speak back to elites, and citizens can communicate amongst each other directly. Where once there were only a limited number potential sources of information, now citizens can choose among an almost infinite variety of sources and interpretations around which to form their beliefs.

In contexts where the mainstream media or government are less established or less trusted, this broadening of the information environment can enhance the credibility of that information from community sources, leading to a greater likelihood that information will lead to updating, particularly among uncertain individuals. Although this enhanced credibility may foster updating of beliefs in the direction of the new information, this updating may not bring opinions closer to accurate understandings of the world as Bayes assumed.[1] For information provided through social media, the "audience" plays a central role in both creating and moderating content, which may introduce individual-level biases into the information environment based on the preferences of the crowd. As a result, although the peer-to-peer nature of social media may be more trusted by individuals in low trust contexts, it may also be more prone to unanticipated error or distortions shaped by the interests (and potential lack of expertise) of those who disseminate it. If information about the world is used to shape our attitudes and political beliefs, this has significant implications for political behavior in communities where the majority of the population primarily gets their information from social media sites.

In the next section, I review the existing literature on attitude formation and social media as a means for the provision of information. I then present my theory, focusing on how the strength of prior beliefs and source credibility in contexts where trust in government and the media are low affects political attitudes. In the final section, I highlight empirical implications of my theory, which I will test in **Chapter 4** and **Chapter 5**.

[1]The word "information" is used in different ways across different strands of research. In theories that involve Bayesian updating and development economics, information is often assumed to be factual. In public opinion literature, the term "information" is used to represent any for of persuasive communication, whether it is truthful or not.

3.1 Past Research on Attitude Formation

Existing literature on how persuasive communications shape political attitudes and actions straddles two separate fields in political science, which rarely interact. The literature in American Politics, influenced by the field of psychology, evaluates how persuasive communications shape political attitudes, with limited differentiation by medium and with a focus on elites. The literature in Comparative Politics, largely shaped by development economics, focuses on how information provided to voters impacts accountability among politicians. This research often fails to explain how the prior beliefs of respondents affect their reception of the provided information. In this section, I discuss the main elements of these existing literatures and highlight their limitations. I then build on these insights to construct a theory of public opinion formation in the social media age.

Public Opinion in American Politics

For decades, scholars of American politics have theorized about the origins of public opinion, and in what contexts perceptions of voters may change or be changed. In one influential branch of this literature, Zaller (1992, pg.,49), argues that citizens express their attitudes through a "top of the head" logic, which requires them to average "across the considerations that are immediately salient or accessible to them." As a result, individuals tend to prioritize recent considerations "and bring them to the top of the head for use" (pg., 48). In this framework, elites can seek to alter public opinion by crafting persuasive messages that are either more salient or easier to recall for voters.

The success of persuasive communications in altering public opinion is, however, mediated by a variety of factors. The credibility of the source (Popkin and Popkin, 1994; Lupia, 1994; Sniderman, Brody and Tetlock, 1993), political awareness (Cobb and Kuklinski, 1997; Druckman and Lupia, 2000; Gaines et al., 2007) and the strength of partisan attitudes (Tormala and Rucker, 2007; Taber and Lodge, 2006; Redlawsk, Civettini and Emmerson,

2010) are cited as factors that contribute to the individual's willingness to incorporate new information into their evaluations, for example, of a politician's trustworthiness.

According to many studies, mass media assists in this process by taking its framing cues from elites (Zaller and Chiu, 1996; Cohen, 1963). Media executives also may set the agenda. For example, Iyengar and Kinder (2010) argue that the attention the media pays to certain issues over others, as well as the ordering of those issues, may mold public opinion. DellaVigna and Kaplan (2007) find that the overt political biases of Fox News boosted Republican turnout in districts where it was first rolled out on television. By contrast, Gerber, Karlan and Bergan (2009) find that it is not the slant of the media, but mere exposure to more information that may matter for belief formation.

To date, this avenue of research has largely focused on the processes by which elites shape public opinion through more curated information mediums such as television or newspapers. The focus of many theorists on elite-driven, top-down persuasion in shaping public opinion may be less applicable for social media, where audiences are encouraged in a bottom-up fashion to "'actively' shape 'media flows'" (Fuchs, 2017). In these settings, elite influence may still matter, but such influence can also be counteracted by public debate, and a viral message from an average user can be just as visible as elite-driven content. For these reasons, existing theories of public opinion formation that focus on the role of elites – particularly through mediums such as television, newspaper, or radio – may fall short.

Information and Accountability in Comparative Politics

In the field of Comparative Politics, scholars who study the role of information in political behavior largely draw on research from development economics. Thus, information is, for the most part, assumed to be accurate. The dissemination of false information and other persuasive messages is seldom considered. The provision of new information is seen as a way for citizens to become better informed about politicians, which should lead to greater accountability. In this framework, information leads to a decrease in information asymmetries

between politician and voters. In theory, voters are then better able to evaluate their elected officials and hold them accountable for their actions.

The basic theoretical foundation for this research stems from the principal-agent problem, first detailed by Eisenhardt (1989). In this framework, an actor, referred to as an agent, takes action on behalf of another actor, known as the principal. Both the principal and the agent have their own interests, which may not align. In deciding whether to act on the principal's preferred action, the agent may weigh their own interests against the costs of fulfilling the principal's objectives. The challenge in many of these scenarios is that the agent is often privy to information that the principal does not have, but needs in order to determine their preferred action. In electoral democracies, incumbent politician should in theory act on the interests of voters – their principals – but they know the costs required to fulfill this political agenda, whereas voters do not. The information and accountability literature in comparative politics seeks to resolve this disconnect by providing citizens with more information about their "agents," which should allow them to close information asymmetries and sanction or reward their efforts.

Although the "information and accountability" literature is voluminous, its findings remain inconclusive. Banerjee et al. (2011) find that providing citizens who live in informal housing with information about incumbent politicians leads to increased turnout, decreased vote buying, and higher quality candidates as compared to slums that do not receive the information intervention. Humphreys and Weinstein (2012) find that information has little effect in changing politicians' behavior or leading to punishment of poor performers. Chong et al. (2014) find that information about incumbent corruption leads not only to lower support for the incumbent but also to a decrease in participation across the board. Lieberman, Posner and Tsai (2014, pg., 70) look at the effect of providing information about education on citizens' public and private behavior. While their findings are inconclusive, they posit a number of reasons as to why: "for information to generate citizen action," they argue "it must be understood; it must cause people to update their prior beliefs in some manner; and

it must speak to an issue that people prioritize and also believe is their responsibility to address."

Building on these insights, Dunning et al. (2019, pg., 6) move the research agenda forward by considering whether the information treatment is considered "good news" or "bad news," based on an individual's prior beliefs. Across seven different studies, they find that information interventions designed to improve government accountability "simply did not induce voters to update their beliefs or produce meaningful change in their perceptions." They also point to a significant issue with past research by cautioning against "using blanket informational interventions to improve democratic accountability." Although several of the studies ask voters to rate their certainty about their priors, they do not exploit the variation in the strength of priors in their main analysis.

Other recent studies in Comparative Politics do exploit this variation. For example, Brierley, Kramon and Ofosu (2018) find that new information provided through political debates does influence voters. Citizens who self-identify as "strong partisans" may be more willing than individuals who identify as "weak partisans" to respond to arguments made on both sides of the political spectrum. In a study related to my proposed research agenda, Arias et al. (2018) shift the focus to the heterogeneous effects of prior beliefs about malfeasance on voter behavior. They find that, after malfeasance revelations, voters who learned illegal activity was higher than their priors, or who updated their beliefs unfavorably were less likely to reward politicians.[1] Although the literature in Comparative Politics has progressed, further efforts to sharpen the measurement of prior beliefs (and the strength of prior belief) and to more fully measure the role of attitude strength in shaping posterior evaluations have been lacking to date, and may help to explain why consensus around this subject remains inconclusive.

[1] Dunning et al. (2019) do interact certainty with their treatment in an appendix analysis and find it has no effect on voter behavior. Both Dunning et al. (2019) and Arias et al. (2018) measure prior certainty by asking voters "How certain are you about your response to this question?" and give them the options "Very certain", "Certain", "Not certain", "Very uncertain".

3.2 Theoretical Contribution

The theory that I set forth in this chapter builds on the theoretical shortcomings of the existing literature. As demonstrated in **Chapter 2**, in contexts outside of the United States citizens who have *less* trust in both the government and elites are *more* likely to turn to social media for information about politics. I argue that two main factors will shape when new information leads to updating: (1) the certainty of prior beliefs; and (2) trust in the sourcing of the information. While the former point is universal across contexts, I argue that the latter is context-specific, varying based on the political environment of those who receive the information.

The Role of Prior Strength in Information Processing

I take as a starting point the literature on attitude formation, persuasion, and public opinion, which is well-established and heavily theorized in the American Politics field. I argue that more information, if it challenges prior beliefs and is salient to voters, may, in theory, lead to a reassessment of prior opinions and catalyze a shift in policy preferences. This logic provides the foundation for Bayesian theory, whereby an individual may come to a situation with a prior idea about the probability of observing some phenomenon in nature, but will update their beliefs about that phenomenon based on new information that they receive (posterior probability) (Gerber and Green, 1999). In a Bayesian framework, an individual's priors, or what they believe, will shape how they respond to new information (Ajzen and Fishbein, 1977). Mathematically, Bayes Theorem can be expressed as follows:

$$f_{posterior}(\theta \mid \text{data}) = \frac{f_{prior}(\text{data} \mid \theta) \cdot f_{prior}(\theta)}{f_{prior}(\text{data})} \tag{3.1}$$

In this calculation, the "data" is the new information and θ is the distribution of beliefs, which the new information may affect. $f(\theta)$ represents the *prior belief* about the world or

how the world works and $f_{prior}(\text{data} \mid \theta)$ represents the *likelihood* of seeing the evidence you actually saw (in the new information), given the current belief. These two factors, when evaluated together, make up the *posterior belief* ($f_{posterior}(\theta \mid \text{data})$), or the updated belief, given the new evidence.

To determine if information will lead to a change in beliefs, we must understand where each individual's current belief about the how the world works ($f_{prior}(\theta)$) stands, and the probability of seeing the evidence conveyed in the new information given your prior understanding of how the world works ($f_{prior}(\text{data} \mid \theta)$). Thus, while an individual can be wrong about their beliefs, it is important to consider how strongly those prior beliefs are held – or how likely the individual is to discount new evidence that is inconsistent with their priors. As such, the degree to which new information will lead to a reevaluation of prior beliefs ($f_{posterior}(\theta \mid \text{data})$) stems from the strength of those priors. In this sense, a person with weakly held priors (50%-50%, or a toss-up) has a near uniform distribution of beliefs about a particular subject, while a person who has strong priors has a belief distribution with a steep peak. In Bayesian updating, very strong priors, or a high degree of certainty about what is true, will lead to minimal updating based on persuasive communications due to the fact that the information is at odds with the individual's strongly held prior beliefs.

American politics scholars have indirectly taken into account prior strength when measuring biased information processing. For example, the idea that citizens seek out information that is preference-consistent may have to do with the strength with which their prior beliefs – particularly their partisan attitudes – are held (Ditto and Lopez, 1992; Zaller, 1992; Taber and Lodge, 2006; Krosnick and Abelson, 1992). More directly, attitude strength has been cited as one explanation for why individuals engage in the process of selective updating, otherwise known to psychologists as motivated reasoning (Brannon, Tagler and Eagly, 2007). In motivated reasoning, individuals may discount, disregard or distort evidence in an effort to make decisions that reinforce their existing beliefs. In this process, what seems like a failure of information to shape attitudes may actually be a deliberate process to avoid the

discomfort of cognitive dissonance when partisan ties are fully developed or, more broadly, prior beliefs are strongly held. According to Festinger (1957), cognitive dissonance arises when beliefs are not congruent with actions or are in competition. Cognitive dissonance will "motivate the person to try to reduce the dissonance and achieve consonance" and to "actively avoid situations and information which would likely increase the dissonance."

Given the potential for prior attitudes to lead to updating failures if the evidence contradicts their current belief, I argue that for information to have any effect on political attitudes or voter behavior it is important not only to understand what an individual's beliefs are, but also to evaluate the strength with which those priors are held. If an individual holds entirely false perceptions, but believes those very strongly, evidence that contradicts those beliefs will produce little updating. By contrast, if a person holds entirely false perceptions, but is less confident in their beliefs, disconfirming evidence may lead them to change their attitudes. This theoretical expectation is highlighted in **Figure 3.1**.

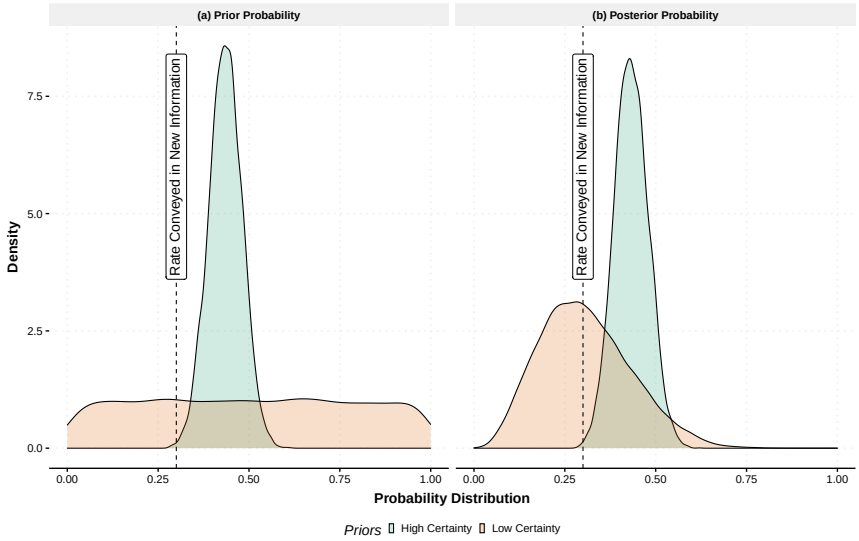

Figure 3.1: The Importance of Prior Strength in Shaping Updating - Figure 3.1 (a) represents the prior beliefs of two hypothetical individuals. Both believe the same thing, but one is more certain about their beliefs, for example of the crime rate, and one is less certainty about their beliefs. Both individuals receive the same information input about the true crime rate in their neighborhood. The information has a greater effect in shifting the posterior probability of the individual with the less precise distribution, as demonstrated in **Figure 3.1 (b).**

Figure 3.1 illustrates how updating occurs based on new information that is inconsistent with prior beliefs for two hypothetical individuals with different levels of certainty about their beliefs. The figure on the left shows the prior probabilities for individuals who are highly certain (green) and uncertain (orange) about their beliefs. The uncertain individual has a wider probability distribution with respect to what could possibly be the true value this hypothetical scenario. By contrast, for the highly certain individual, there are fewer alternatives that could potentially be true.

The figure on the right illustrates how these probability distributions shift based on new information that is inconsistent with the prior beliefs of these individuals. In the posterior probability density curves, the individual with the more uniform density curve distribution moves closer toward the rate conveyed in the new information, and their probability distribution becomes less uniform. The high certainty individual does not change their beliefs, but may become less certain about them (as demonstrated in the above Figure) or may discount that information completely and not shift their beliefs at all. In general, shifts in the posterior probability distribution may become more or less sharp depending on whether or not the new information is consistent with prior beliefs.

Bayesian updating does not consider the source of new information. However, in the real world, information does not simply appear in the real world as fact. Instead, it is tied to a source which provides that information. This source is not a neutral actor, and instead carries with it preexisting, *external* connotations or perceptions unique to the context in which the information circulates. As such, while information processing may follow a mathematical formula, the perceived *credibility* of that information matters with respect to updating. If a source is viewed as lacking in credibility, no amount of information will lead to changes in prior perceptions. This the credibility of the source also factors into the likelihood of seeing the evidence, given the prior beliefs people hold.

Source Trust and Updating

In addition to the theoretical expectation that individuals who are less certain about their beliefs will be more likely to change their perceptions, I also argue that the source of that information matters. Much of the research highlighted above has assessed the role of elite-driven persuasive communications via television, newspapers or radio in altering public opinion or provided neutral content to study participants.

In Bayesian theory, the credibility of the information source is not considered at all. Despite this omission, some research has already complicated Bayesian theory to incorporate features that make people more or less receptive to new information. For example, people are more likely to accept new information from those they recognize as experts (McGuire, 1985) or from co-partisan sources (Zaller, 1992). These theories of opinion change were developed to explain the effect of one-way communication from elites to ordinary people, and they lead to clear-cut expectations about how exposure to new information in the media influences the opinions of those exposed. When elites share particular values, the citizens who are most exposed to elite discourse are most likely to adopt the same values. When, however, elites disagree, then citizens tend to adopt the views of those they find more trustworthy (co-partisans in the U.S.) and reject the information provided by those they view as less trustworthy. Regardless of whether elites are unified or divided, this process tends to result in popular views biased in the same direction as those of elites.

Since the bulk of this research was conducted, technological change has introduced new factors that shape the information environment, which in certain contexts, might affect the receptiveness of individuals to new information. In the field of political science, these factors have yet to be incorporated into theoretical expectations. Yet, scholars – primarily in the field of communications – have recognized and begun to explore how information provided through the internet might be different from information provided through other mediums. Agre (2002) argues that unlike other types of technology, the Internet has the potential to

foster social networks and "bind people together into a polity." In a similar vein, O'reilly (2005) argues that social media platforms differentiate themselves because they are built around the "architecture of participation" and remove bottlenecks to spreading information, such as media executives.

Although social media share many attributes with other information mediums, they differ in at least two important ways: low barriers to entry and reliance on user-generated content. Low entry barriers make the gatekeeping of the spread of political information much less effective by leveling the playing field between traditional providers of information and the masses (Zhuravskaya, Petrova and Enikolopov, 2020; Farrell, 2012). The reliance on user-generated content (and user-driven policing of that content) means that information filters and information demand operate in a feedback loop shaped by and for everyday users. Media curation of content by professionals is diminished and instead replaced by the "wisdom of the crowd."

This deluge of information from people "like you" does not exist with other information mediums, and I argue that it has two distinct effects on the information environment: (1) people in contexts with historically low government or media trust may view this content as more credible, increasing the likelihood that they update their beliefs in response; and (2) it is likely to be prone to exaggerations and exclusions based on the preferences or perceptions of different networked communities, meaning that what people view as more credible may also be biased in unpredictable or unanticipated ways.

In contexts where government or mainstream media are less established or seen as biased, the distinct attributes of social media may be viewed more favorably, leading to greater information processing. This may be why in contexts outside of the United States, citizens who have *less* trust in the government and elites are *more* likely to turn to social media for information about politics, as demonstrated in **Chapter 2**. The removal of bottlenecks and information curators allows "new entrants" to compete on an equal playing field alongside the establishment – a reality which may be preferable to people living in countries where

trust in government remains low, often due to past histories of authoritarianism or civil conflict; or where media is less established, less independent, or was previously coopted by undemocratic processes.

"Crowd-sourcing," a term coined in 2006 to describe a phenomenon whereby the masses can harness their collective knowledge to solve problems over the Internet, becomes a more reliable and more trusted source (Howe, 2006). The idea of the collective "wisdom of the crowd" was initially tapped into by technology companies, but today it extends toward all avenues of civic life, ranging from community building, monitoring and mobilization, to documenting noteworthy incidents in real-time (Surowiecki, 2004). In his book *Here Comes Everybody*, Shirky (2008, pp. 55) calls this the "mass amateurization of efforts previously reserved for media professionals." In these contexts, "everyone is a media outlet," which "undoes the limitations inherent in having a small number of traditional press outlets."

Combining these two expectations, I expect individuals with strongly held priors to be more resistant to new information. Likewise, I expect individuals to be more open to updating their beliefs, if these priors are weakly held. While this may be true across all different types of persuasive communications, in contexts of low government or media trust, I expect exposure to information on social media to make these effects more salient than does the information provided through other channels. This is due to the fact that the information comes from avenues perceived to be more trustworthy. Both the strength of prior beliefs and source trust will lead an individual to accept or discount that information, which will then trigger a shift in public opinion and policy preferences. Drawing on Zaller (1992), these theoretical expectations are highlighted in **Figure 3.2.**

In this Figure, the blue line represents an individual with weak priors and the red line represents an individual with strong priors. The x-axis represents how much information an individual receives from social media and the y-axis represents how much that information affects their policy preferences. Moving across the x-axis, I argue that in both cases, the more the exposure to social media, the more likely an individual is to update their beliefs. This is

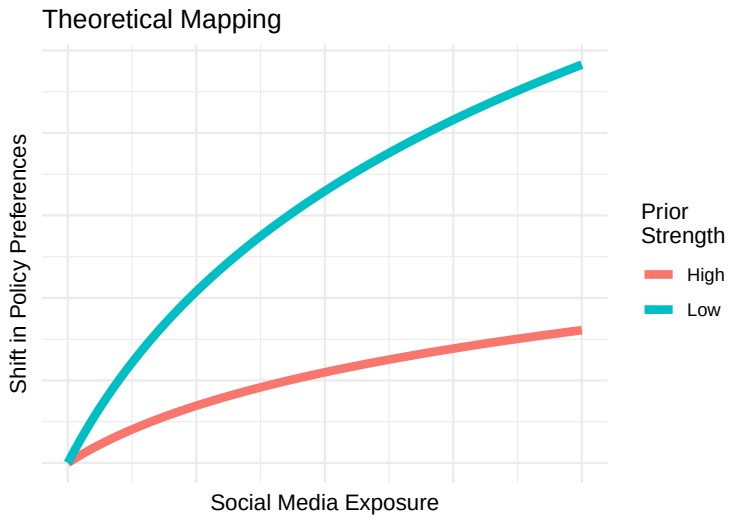

Figure 3.2: Theoretical Expectations - Figure provides an overview of broad theory linking prior strength to social media. I argue that social media will be more salient in changing public opinion due to the low-barrier to entry and reliance on user-generated content. However, this effect is mediated by individuals who hold strong priors about crime in their communities. In these cases, the effect of information will be less effective.

due to the perceived credibility of the source, which in settings of low trust in government and media is viewed more favorably, as a result of high historic distrust in traditional purveyors of information. At the same time, this effect is mediated by the strength of an individual's prior beliefs. Individuals with strong priors who receive the majority of their information from social media will update much less than individuals with weak priors, due to their receptiveness to receiving new information that may challenge their priors.

Put a different way, we can think of two factors – strength of beliefs (internal) and source credibility (external) – as variables that when combined together affect the "potency" of the information on updating. In specific contexts where trust in government and media is low, we can expect the greatest amount of updating to occur among individuals who are uncertain and acquire information from social media. We can expect a small amount of updating to occur among individuals who are *either* uncertain *or* receive their information from social media. Finally, we can expect no updating to occur among respondents who are both highly certain about their beliefs *and* obtain their information from traditional channels. These expectations are detailed in **Figure 3.3**, which distills these expectations into a 2x2 matrix.

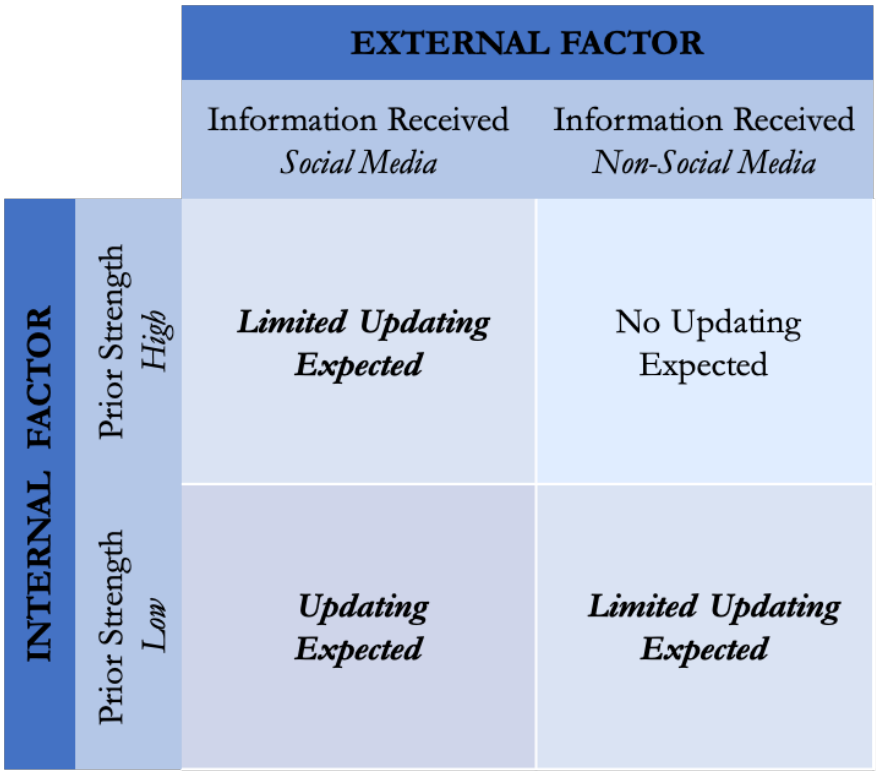

Figure 3.3: 2x2 of Theoretical Expectations - Figure details the expected effect of information, based on the prior beliefs of the recipient (internal) and credibility of the information (external).

Updating Based on Biased Information

In a number of developing country contexts, the power of the crowd has been unleashed to address a wide range of challenges related to disaster relief (Gao, Barbier and Goolsby, 2011), electoral violence (Meier, 2011), and violent conflict (Dowd et al., 2020). Harnessing the collective power of the crowd means that citizens can provide more information, rapidly and at a level of geographic disaggregation that previously may not have been possible, due to resource constraints on the traditional media or government bureaucrats (Bott, Gigler and Young, 2014). Knowing that this information comes from a non-dominant or non-monolithic source may enhance its credibility in the eyes of those who receive it. By contrast, more traditional information pathways may be entirely discredited, and no amount of information will change this perception, particularly among individuals with strongly held beliefs.

Although I argue that in low-trust contexts, citizens are more likely to trust information (and subsequently change their beliefs) when it comes from people "like them," relying on the crowd for this content also has its own downside. Although "mass amarteurization" may produce a greater volume of content on a far wider range of topics, the information that citizens gather will be filtered through their own prior perceptions of what they view as "newsworthy," rather than the more uniform filters determined by a media executive or professional (Shirky, 2008). The filters tied to more traditional information purveyors were prone to their own implicit biases about what was worthy of being shared, but they were uniform within an outlet, and usually more similar than not across different media providers (Nye, Cherny and Alterman, 1990). Thus, the biases were at least implicitly known and understood. Furthermore, traditional media outlets must – at least in some respects – guard against reputational risks tied to low-quality or highly skewed content (Cagé, 2020). With the crowd in charge of what does and does not get shared, information filters may be vastly different, reputational risks largely eliminated, and preferences for sharing specific types of information may be distinct in specific contexts.

These biases can be both broadly applicable across all content creators or specific to certain contexts and individuals. Broadly speaking, people have a tendency to gravitate toward content that is more shocking or fear-inducing. This is a common feature of human psychology, whereby citizens are more likely to engage with negative information due to what is known as a "negativity bias" (Ito et al., 1998). This is why traditional media sources often pay disproportionate attention to bad news as a function of their business model (Altheide, 1997) and why political campaigns can be more effective if they focus their messaging on topics that evoke fear or anxiety among potential voters (Ansolabehere and Iyengar, 1995). In these contexts, negative sentiments are more memorable and "contagious" than positive emotions (Rozin and Royzman, 2001; Soroka, 2006).

For information distributed on social media, the potential to focus on content that evokes fear and anxiety is a natural extension of this preference. Without journalistic oversight and due to the vast amount of content that can be generated, social media has the potential to further "unleash and intensify collective alarm" (Walsh, 2020). This content can lead digital platforms to become "instruments of panic production" and greatly exaggerate any potential real-world threat (Bebbington et al., 2017). This may be particularly true in in more affluent communities, or in communities that are less exposed to violence due to cognitive biases, as detailed in Kahneman and Tversky (1979). In contexts where violence might be a rare event, people will either ignore or more heavily weight improbable incidents and exaggerate their importance far more than they would if they thought in terms of probabilities. Kahneman and Tversky (1979) note, "because people are limited in their ability to comprehend and evaluate extreme probabilities, highly unlikely events are either neglected or overweighted, and the difference between high probability and certainty is either neglected or exaggerated." (p. 283). Thus, a single shocking incident may be overweighted simply because it is probabilistically unlikely. This phenomenon is related to what Tversky and Kahneman (1974) refer to as a *salience bias*. The salience bias occurs when there is an unexpected contrast between an event or object and its surroundings – for example, a

black sheep in a flock of white sheep. This leads people to focus more on the black sheep, simply because, in this case, it is more prominent, even though it is far rarer. The same type of phenomenon can occur with events – for example, violence in a normally peaceful area, which leads people to focus more on the single incident of violence, because it is emotionally evocative, than on the much more frequent days without any incident.

By contrast, in poorer communities a negativity bias might still shape what people respond to and share, but violent events will be less salient because they are probabilistically more common. Furthermore, although social media content produced by people living in less well-off, violent areas also contains plenty of bad news, many communities have taken to social media in an effort to explicitly counteract a common narrative of these places as rife with poverty, violence and/or hardship. This type of content, which is often the sole focus of traditional media coverage, is so common that it has been popularly labeled "poverty porn" (Feltwell et al., 2017). For example, Moors (2019) documents how after the Flint Water Crisis in 2016, citizens in the community took to social media to promote "#FlintFwd" as a way to challenge the dominant narrative. Their objective was to counter negative coverage of Flint and share what they viewed as the real Flint, with all of its "vibrancy and dynamism" (pg., 814). The same phenomenon has occurred in Rio de Janeiro, where younger people in Rio's poorer favela communities have taken to the Internet to share an alternative narrative of their daily lives, one that counters the perceptions of favelas as existing in a perpetual state of quasi-civil war (CESeC, 2019).

These cognitive biases undoubtedly shape the type of content that is and is not shared across different communities on social media. As a result, if citizens in low-government and media trust settings are more likely to update their beliefs based on social media content, this variation in content may have serious implications for political belief formation and perceptions about the external environment.

3.3 Empirical Implications

My main theoretical expectation is that as citizens rely more on social media for information about their environment, they will be more likely to change their beliefs, particularly if those beliefs are not strongly held. This over-reliance on social media, however, may introduce distortion into the information environment, rooted in the cognitive biases of the individuals generating content. The main observable implications from my theory are detailed below and tested in subsequent chapters:

- **H1).** Information provided through social media will alter respondents' concern (or worry) about crime more than information from the government, with the effects concentrated among respondents who were initially uncertain about their knowledge.

- **H2).** Information provided through social media will be more likely than information from the government to lead to changes in policy preferences, with the effects concentrated among respondents who were initially uncertain about their knowledge.

- **H3).** Respondents are more likely to trust information from social media than information conveyed through traditional media due to attributes inherent to these platforms.

-H4). Social media content will be more likely to cover shocking or fear-inducing content than mainstream media.

-H5). Citizens are more likely to engage with shocking or fear-inducing content from social media than mainstream media.

-H6). Shocking or fear-inducing content on social media will be more common in places where citizens have higher uncertainty about the frequency of these events.

-H7). Citizens will be more likely to engage with shocking or fear-inducing content on social media in places where citizens have higher uncertainty about the frequency of these events.

I address hypotheses 1-3 through a series of survey experiments in **Chapter 4**. I address hypotheses 4-7 in **Chapter 5** using data from Facebook community pages and government violence statistics at the neighborhood level. More detail on the research design is provided below.

3.4 Research Design

To study the transformation of public opinion formation in the social media age, I utilize a mixed-method research strategy that combines observational data and survey experiments. To examine the remaining aspects of the theory, I rely on survey experiments and large-N text analysis. I utilize both information experiment questions and conjoint questions to identify a causal effect for specific hypotheses posed by the theory. In the information manipulation, I assess how perceptions of insecurity and support for increased policing shift when the source of crime information comes from social media versus when it comes from more traditional platforms such as the government. I also evaluate the relationship between the strength of prior beliefs about crime, and the magnitude of the change. I measure the strength of prior beliefs by asking respondents five knowledge questions, including the degree of certainty with respect to those questions. The ratings of their certainty are averaged together to create an "uncertainty index." In the conjoint question, I assess which aspects of the information provided make it more trustworthy to citizens, varying the range of categories to include: (1) the social networking dimension (e.g., likes and comments); (2) the sourcing of information (e.g., your community, your friends, the media, or the government); (3) the volume of the information (e.g., multiple short reports or one larger story); and (4) the presence of accompanying materials (e.g., photos/videos). This survey of 1,245 citizens was conducted in December 2020 in collaboration with Quaest Pesquisa e Consultoria, a Brazil-based polling firm. More details on the survey experiment as well as results are included in **Chapter 4**.

In **Chapter 5**, I draw upon vast amounts of Facebook data, both from official news sources and community Facebook pages that are citizen run and include citizen generated content. In total, I examine the content and reactions (comments, likes, shares, and positive and negative reactions) to nearly two million Facebook posts between 2013 and 2019 from

21 popular news sites and 221 community pages.[1]. Using regression analyses, I examine the relationship between posts about crime, reactions to these posts and sociodemographic features, such as the Social Development Index (IDS) of each neighborhood, of where these posts are being shared. I also look at overarching patterns in regular news media, which is curated by media elites, as compared to the more crowdsourced variety shared in community pages.

Although none of these pieces alone provides full support for the theory, collectively, each offers empirical evidence in support of the theory I have laid out in the previous chapters. Although this theory is a general theory of change, with no specific geographic reference point in mind, I conduct all of my empirical research in Brazil, with a primary focus on the city of Rio de Janeiro. In addition, I specifically focus in on information about crime. In theory, I would expect any salient information where interest is widespread but exposure is not necessarily uniform across space to yield similar results (e.g., refugee resettlement or immigration).

3.5 Conclusion

The theory described above is broadly applicable across contexts where trust in government and media are less well established. With respect to the information being shared, the theory relates more generally to information that is not universally known across geographic space. I am specifically interested in an issue that is politically salient, about which knowledge may be unevenly distributed but interest is widespread. In a country such as Brazil, there is a natural focus on violent crime. In Brazil, crime consistently ranks as **the** top political issue among voters. This is particularly true in urban areas, where violence often coexists alongside extreme inequality.

[1]Popular news sites are taken from the Reuters Digital News Report Brazil country report, which provides a list of the most popular media outlets annually by country. More detail on the identification and inclusion of community pages is discussed in **Chapter 5**.

The organic rise of online, citizen-run groups that provide information on crime is one strategy used to address insecurity. It mirrors the OECD (2009) recommendation to adopt "innovative information communication technologies for data collection, communication, visualisation and analysis" to address security threats. However, the effect of this rapid, real-time, citizen-sourced data remains underexplored, with most academic research in other geographic contexts focused on how representative the data is of actual on-the-ground events.[1] Although the use of this type of technology to fight crime is a global phenomenon, in low-income countries where violence is a constant in some neighborhoods, this real-time knowledge may be vital to citizen security. However, in areas where crime is lower, information about one criminal act may exert outsized influence on the imagination of voters. Thus, the impact of this type of communication network may produce a broad range of political outcomes and differentially alter citizens' political attitudes and perceptions of their own communities, based on what is shared.

[1]For other examples of the crowd-sourcing phenomenon, see: Unver and Alasaad (2016); Mäkinen and Wangu Kuira (2008); Israni, Erete and Smith (2017); Dowd et al. (2020); Roberts and Marchais (2017); Erete (2015); Van der Windt and Humphreys (2016); Pridmore et al. (2019).

4

Trust Thy Neighbor: How Social Media Shapes Political Attitudes about Violence

"Meaning just ain't in the head."

– *Hillary Putnam, The Meaning of "Meaning"*, 1975

In 2013 a video from northern Mexico depicting the decapitation of a young woman, presumably at the hands of a narcotrafficker, went viral across social media. The video made waves not only because of its shockingly gruesome portrayal of violence, but also because mainstream media largely ignored it. Prior to the viral video, the Mexican government had ignored the event and referred to the panic surrounding gang violence as a "collective psychosis." The video provided clear evidence to refute this claim. In Brazil, residents in several favelas have begun to use social media to organize "collectives," such as *Coletivo Papo Reto* in *Complexo Alemão*, to provide information to the community about a variety of subjects, ranging from neighborhood meetings to police brutality. Their activities are so widespread and their networks so vast that last summer, during a military operation in one

of Rio's poorest favelas, police arrested individuals who belonged to social media groups in order to prohibit them from sharing images and videos online.[1] These two cases are part of a broader pattern whereby social media has begun to disrupt the historic flow of information about violence from government to citizens, but the question remains: how transformative is this change?

An influential body of literature in political science examines the effect of information in shaping public opinion and voter behavior. Although this research spans many different countries and political systems, to date much of this work either evaluates the role of information provided through television, newspapers, and radio, or fails to delineate between modes of information provision at all. Yet in many countries, such as Brazil, social media platforms are increasingly a priority source for acquiring information about the world. Despite this clear trend, there is limited study of the ways in which information obtained through social media may be different from other types of sources in driving public opinion and voter perceptions in these contexts.[2] In particular, the low barrier to entry and the reliance on user-generated content characteristic of these platforms distinguish them from other information mediums. These distinctions may have implications for information processing and attitude formation, particularly in settings where trust in government and political institutions is low.

The research presented in this chapter seeks to fill a significant theoretical gap by exploring the effects of community-driven information provided through social media – particularly information about salient topics, such as crime – on political attitudes and perceptions of insecurity among voters. Specifically, I ask: (1) are respondents more likely to change their beliefs about violence, if the information comes from people like them? (2) what role does the strength of previous beliefs play in this calculation? and (3) if beliefs are more malleable

[1]Information provided during interviews in August 2018, when the operation in *Complexo Alemão* occurred.

[2]In one of just a few empirical studies on the subject more broadly, Guess, Nagler and Tucker (2019) found that the fakes news on social media may be less influential than imagined, and may only affect certain types of people – namely, older voters.

when information comes from social media, what about these platforms might make this content more credible to respondents?

Uniting insights from American and Comparative Politics, this chapter advances our understanding of public opinion and voter behavior in the social media age by evaluating the conditions under which social media may matter for political belief transformation. From a theoretical perspective, I argue that in a Bayesian framework, prior beliefs and the strength of those beliefs mediate how and when information leads to updating. This is true across all different types of information processing. However, I also contend that the processes by which information provided through social media may shift attitudes are distinct from other types of media due to the "many-to-many" nature of these communications platforms. In these types of communications, many people can create content, but many people create content, which many people can receive and respond to. Social media information, due to the nature of its sourcing, may be especially important in countries where people distrust traditional purveyors of information, including the media and the government. The fact that social media information comes from "people like you" may make it more credible and therefore lead people to be more willing to update their beliefs about the topic of that information.

I study the effect of this shift in the information environment on public opinion formation in Rio de Janeiro, Brazil, a context that both benefits from and is hindered by widespread information transmission through social media. In particular, I focus on information about crime, a politically charged issue across the city, where exposure to violence varies dramatically by neighborhood. To measure the strength of prior beliefs, I developed an index ranging from 0 to 1 that is based on respondents' self-reported uncertainty with respect to five questions about violent crime. I then randomized participants to receive one of two information prompts about neighborhood violence – one included standard information from the government and the other included crowd-sourced information from a social media platform.

I find that respondents are more likely to update their beliefs about violence in their

community and their support for increased policing when they receive information from a social media source. However, I find null main effects due to the fact that these effects push in different directions, depending on the certainty level of respondents. While highly certain respondents are significantly more likely to double down on their beliefs, less certain respondents are significantly more likely to reverse course. I observe no such heterogeneous effects among respondents who receive information about violence from the government.

Understanding these heterogeneous effects is beyond the scope of this dissertation, but these findings suggest that while uncertainty may play less of a role than expected, the nature of the information does matter. To further explore these results, I conducted a conjoint survey experiment, where I showed respondents two pieces of social media information (one related to violence and one to government services) four times each. I then varied the attributes tied to this information on a range of levels: (1) source, ranging from experts to community members; (2) speed, meaning the time between events and information delivery; (3) post engagement in the form of likes, comments, etc., evaluated here as a signal of community-driven support for the message; and (4) evidence associated with new information, such as videos or photos. The main outcome for this conjoint experiment asked respondents "which profile would you trust more."

I find that across the board, respondents are more likely to trust information that does not come from the government, that is provided in real time, that has higher levels of community-level support, and that includes additional evidence. The features are attributes commonly associated with social media information. These findings are consistent regardless of the uncertainty levels of respondents and regardless of whether the information is about violence or the provision of government services. This suggests that information provided through social media may be better suited toward generating trust in certain political contexts., which may shape citizens' willingness to update their beliefs in response.

This chapter proceeds as follows. In **Section 4.1**, I provide a brief review of the theory, detailed in **Chapter 3**, as it relates to the observable implications tested in this chapter.

In **Section 4.2**, I detail the data collection process and how I coded my dependent and independent variables, as well as the research design used for the analysis, presented in **Section 4.3**. In **Section 4.4**, I highlight several implications from these findings and point to possible mechanisms driving these results. In **Section 4.5**, I discuss additional areas for future research.

4.1 Theory Overview

In this section, I provide a brief overview of the theory detailed in **Chapter 3** as it relates to the empirical portions of this chapter. I argue that for information to matter – and for it to be meaningfully incorporated into people's beliefs – it has to have two components: (1) people have to be open to receiving new information; and (2) they have to trust the source of that information.

The idea that new information leads to changes in prior attitudes – a foundation of Bayesian theory – is a common starting point for research focused on public opinion. In a Bayesian world, new information interacts with prior beliefs to shape posterior attitudes. In much of the existing research on attitude formation, this straightforward process is frequently just the starting point. What also matters is how the information comports with an individual's prior worldview (i.e., if it is disconfirming, it may be ignored altogether) or if it comes from a source that is viewed as trustworthy (i.e., someone considered an expert in the area).

To further complicate this dynamic, I argue that for new information to have an effect on shaping beliefs an individual must be open and receptive to that new information. Specifically, an individual may come to new information about a topic for which they have already assigned a specific belief, yet the new information may also challenge their preexisting beliefs. If those beliefs are not strongly held, the new information will likely be incorporated into their updated beliefs, even if it is disconfirming. If those beliefs are held with a high degree

of certainty (even if they are wrong), new information about a given topic will have little to no effect in altering their beliefs. Thus, I argue that the strength of prior beliefs shapes when and if information leads to changes in attitudes or beliefs.

Drawing on research in political communications, I argue that social media has opened up new pathways that citizens can utilize in order to shape their beliefs. Prior research on how information shapes attitudes and beliefs about politics largely focuses on an era of political communications before the advent of social media. Social media, by definition, has transformed the communications environment from a one-to-many dynamic into a many-to-many conversation, whereby citizens can communicate with politicians and people like themselves with the simple click of a button. This transformation of the information environment from a monologue into a multi-way conversation may be particularly useful in contexts of low-or less-established government trust, where the crowd may be viewed as more trustworthy than traditional gatekeepers of information.

I test the following hypothesis in this chapter, which are derived from my broader theory summarized above:

H1) Information provided through social media will alter respondents' concern (or worry) about crime more than information from the government, with effects concentrated among respondents who are uncertain about their knowledge.

H2) Information provided through social media will be more likely than information from the government to lead to changes in policy preferences, with effects concentrated among respondents who are uncertain about their knowledge.

H3) Respondents are more likely to trust information from social media due to attributes inherent to these platforms.

4.2 Data and Research Design

I explore these dynamics through a municipality-representative survey of Rio de Janeiro, in Rio de Janeiro State, Brazil. In 2014 and 2016, Rio was host to two major mega-sporting events, the World Cup and the Olympics respectively, which catalyzed the widespread adoption of social media in response to information holes left by the state government. In particular, this content focused primarily on crime and violence, which remain a huge challenge for the government and major sources of insecurity for citizens across the municipality, despite the implementation of government policies designed to reduce crime and violence (e.g., the Pacifying Police Forces). Although total crime-related deaths have gone down over time, some incidences of violence – such as police killings – are on the rise (**Figure 4.1**). In a municipality-wide survey, nearly three quarters of all respondents either "somewhat" or "strongly" worry about themselves or a family member getting caught in crossfire between the police and narcotraffickers (**Figure 4.2.**) This demonstrates the salience of insecurity in Rio, particularly at the hands of the security sector. In response to these concerns about violence, social media apps in this context are designed to report on violent incidents, hold the security sector accountable for police brutality, and fill gaps in data about violence left by the state.

To explore these dynamics, I fielded a municipality-representative survey of 1,245 respondents in partnership with Quaest Pesquisa e Consultoria, a Brazil-based firm that specializes in public opinion research using social science methodology. The survey ran from November 26, 2020 to January 6, 2021. Quaest recruited participants for this survey from a nationally-representative panel of respondents. Respondents already had been selected from household samples and had indicated that they would be willing to participate in new research in the future. The panel contains approximately 150 thousand people nationwide, with eighty thousand in the city of Rio de Janeiro alone. A stratified random sample from this population was drawn to compose the final sample of respondents, who were then contacted via a mobile

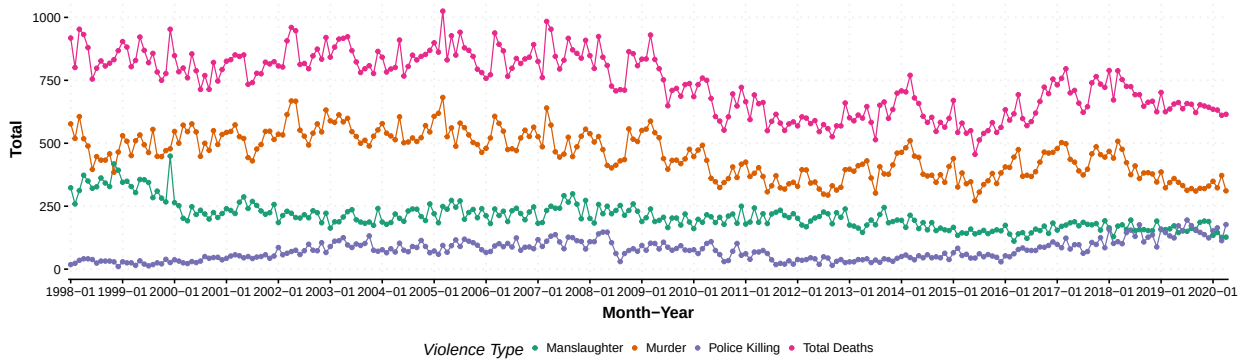

Figure 4.1: Trends in crime-related deaths, by type and overall. - Time series data comes from Rio de Janeiro's Public Security Institute (ISP) and is collected by the Civil Police. The line for "Total Deaths" represents the sum of "Manslaughter", "Murder" and "Police Killings."

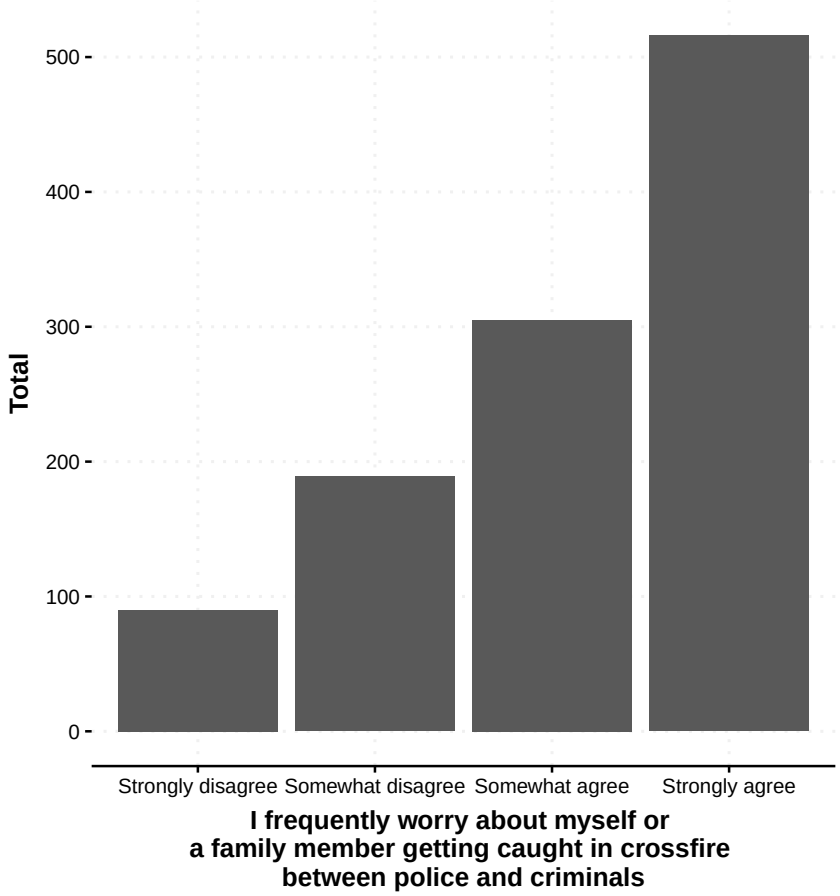

Figure 4.2: Concerns about getting caught in crossfire. - Distribution of survey responses to the question "I frequently worry about myself or a family member getting caught in crossfire between police and criminals" from my survey, conducted in 2020.

app to gauge their willingness to participate in the survey. The final sample is representative of the population in terms of both sex and age, two variables that are both highly correlated with violence and social media exposure. In terms of socioeconomic indicators, the sample skews a bit wealthier than the population average, but given that the research focuses on "online populations," this is to be expected and less of a concern.

To participate in the survey, prospective respondents received a link through an internal Quaest system. The details about how to complete the survey online were then provided through email and text message. I programmed the survey and randomization into treatment conditions online through Qualtrics. Once participants completed the online survey, they received an award of their choice through the Quaest system: they could make a donation; receive minutes on their cell phone; get "Livelo" points to save up for vacations, products, or services; or get "Dotz" rewards to spend online. Rewards ranged in value from R$1 to R$6, depending on their selection, or around US$0.20 to US$1.20.

The questionnaire covered a range of topics, from basic demographics to media habits and experiences with violence, as well as knowledge questions, and two survey experiments, both of which are discussed further in this chapter.[1] In total, respondents answered 41 questions tailored to their neighborhood, including four questions as part of the information experiment and ten questions tied to measuring the strength of prior beliefs, five of which are included in the prior beliefs index described below. These two groups of questions form the basis for the analysis presented below. I discuss each in more detail in subsequent sections. The remaining questions are used throughout the dissertation to provide suggestive evidence to support broader theoretical claims.

[1]The survey instrument and hypotheses for this chapter, and the dissertation more broadly, were pre-registered online: https://egap.org/registry/.

Independent Variable: Measuring the Strength of Prior Beliefs

The main independent variable for the hypotheses detailed above is **the strength of prior beliefs**. To measure this, I ask respondents five knowledge questions. I followed each knowledge question with a question that measure the strength of their beliefs. From a theoretical perspective, I am more interested in how *strong* their perceptions are as opposed to whether or not their beliefs are accurate. This is due to the fact that a person could hold inaccurate beliefs but believe them very strongly, meaning that they are more stubborn in their attitudes and less likely to change their mind regardless of the validity of their perceptions. The five knowledge questions I asked are as follows:

> **Knowledge Question #1:** Approximately how many homicides were there in your neighborhood last year?
>
> **Knowledge Question #2:** On which street is the military police battalion in your neighborhood located?
>
> **Knowledge Question #3:** Has violence gone down in your neighborhood?
>
> **Knowledge Question #4:** Which police force is responsible for investigating crime?
>
> **Knowledge Question #5:** What is the general telephone number of the Military Police?

These knowledge questions were prescreened on Rio de Janeiro residents to ensure that they were at least possible to answer. Some questions were more difficult than others, and questions were ordered from most difficult to least difficult. For each of five knowledge questions in the survey, I asked respondents how certain they were about their answer to the question. In each of these five "certainty questions," respondents who reported being "very uncertain" were coded as a "4"; "somewhat uncertain" as a "3"; "somewhat certain" a "2", and "very certain" a "1." I then calculated the sum of these individual scores and divided

this total by 20 (the maximum level of uncertainty). Higher values correspond to greater uncertainty, whereas lower values correspond to less uncertainty. The final uncertainty index ranges from 0 to 1 based on respondents' self-reported uncertainty with respect to five questions about violence and crime. **Figure 4.3** plots the breakdown of this uncertainty score.

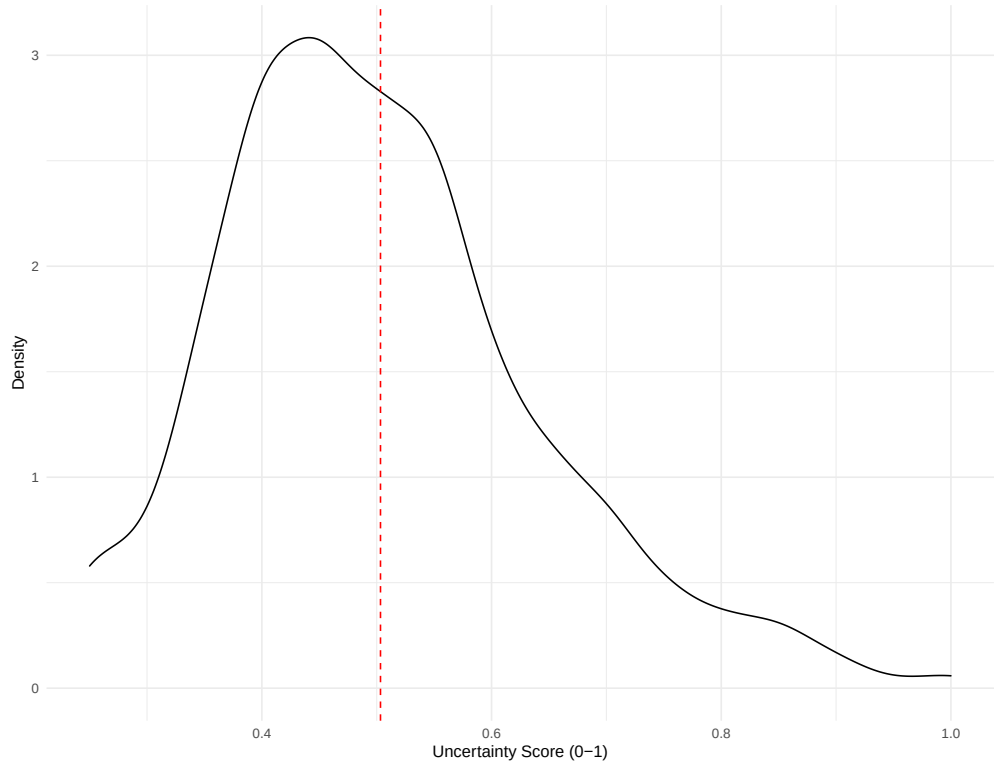

Figure 4.3: Uncertainty Scores Distribution. - Average uncertainty score denoted by dashed red line (0.5033). High values represent high uncertainty.

The average uncertainty score is 0.5033, meaning that respondents hold at least some level of confidence in their responses to questions about crime and violence. 759 respondents are more uncertain in their responses than the average respondent, and 487 respondents are less uncertain than the average respondent. The median uncertainty score is 0.5.

It is important to emphasize that this is an imperfect proxy measure for strength of prior beliefs. Ideally, I would measure certainty around subjectively held beliefs as opposed to more objective questions; however, it is more difficult to validate the subjective beliefs individuals hold. This task becomes possible with factual questions. For the purpose of

this analysis, I argue that individuals who view themselves as *more* certain about factual questions related to violence will be *more* aware of crime and violence as a subject more broadly. I then assume that individuals who hold strong beliefs about these questions will likely also hold strong beliefs about less measurably attitudinal questions related to violence, due to their increased awareness of it as an issue.

Treatment Assignment

I then used a simple random assignment to randomize respondents to receive one of the following pieces of information, which provides **accurate** data about violence in their neighborhood based on their response to a question about what community they live in. The two information prompts are provided below in **Figures 4.4**. Respondents were randomly exposed to only one piece of information, which included violence data from the source provided (either ISP, a government source, or Fogo Cruzado, a social-media, crowd-sourced, social media platform). I explicitly varied the source of the information and the timing that they the information was reported. Due to the nature of the information provided by each source, I also varied whether the incident was a "violent incident" or a "homicide." This is due to the fact that Fogo Cruzado does not always report gunshot violence tied to homicides. However, I do not expect this to affect my results in any major way. If anything, it would dampen the effect of the social media treatment due to the perceived severity of a "violent incident" as compared to the certain severity of a homicide.

Given that the two prompts are providing truthful information at the neighborhood level, it is important to note that respondents might have received *different* information both due to variation across neighborhoods and variation by source. Taking into account this variation, I controlled for the numeric value of information that each survey respondent saw in all subsequent analyses. Additionally, I note that both sources of information are highly correlated, as demonstrated in **Figure 4.5.**

In expectation, random assignment "implies that the observed and unobserved factors

Figure 4.4: Information Experiment Prompt

(a) Social Media Source Prompt Text: According to a local social media site, which provides citizen-to-citizen information on neighborhood gunshots, residents reported [NUMBER] violent incidents in your neighborhood of [NEIGHBORHOOD] between January and August of this year. These incidents were reported in real-time, while the violence was occurring.

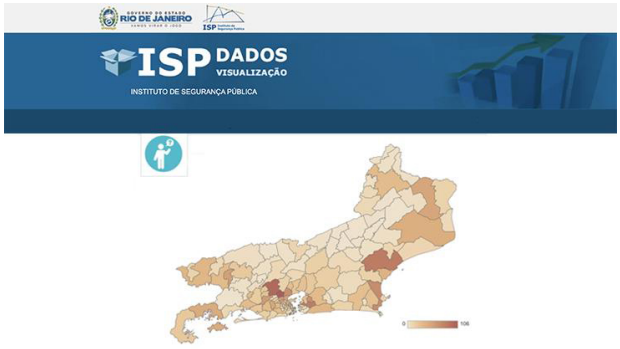

(b) Government Source Prompt Text: According to the Institute for Public Security, which provides government data about violence in Rio de Janeiro, the civil police reported [NUMBER] homicides in your neighborhood of [NEIGHBORHOOD] between January and August of this year. These incidents were reported online a month later.

Figure 4.4 provides an overview of the two information prompts, social media source (a) or government source (b), of which respondents saw one.

that affect outcomes are equally likely to be present in the treatment and control groups" (Gerber and Green, 2012). However, in practice, this may not be the case due to random chance. As a result, it is useful to block on potentially critical variables prior to the treatment assignment, to ensure balance across treatment and control groups. Ideally, I would have blocked the treatment assignment on levels of uncertainty given its central role to my theory and hypotheses. However, the post-hoc need to calculate the uncertainty score using the formula described above made this difficult to accomplish in the same survey period. I did, however, check for balance to ensure that the randomization was conducted properly. I checked for balance on all covariates collectively via an f-test. The f-statistic for the joint hypothesis test is 2.3355 with a p-value of 0.1267. This means that the model, made up of ten different covariates, has no statistically significant power in explaining treatment assignment. I also checked for balance on each covariate in **Table 4.1** by conducting a series of t-tests and comparing the means for treatment and control groups by variable.

Fortunately, there appear to be no major imbalances on most covariates, including the

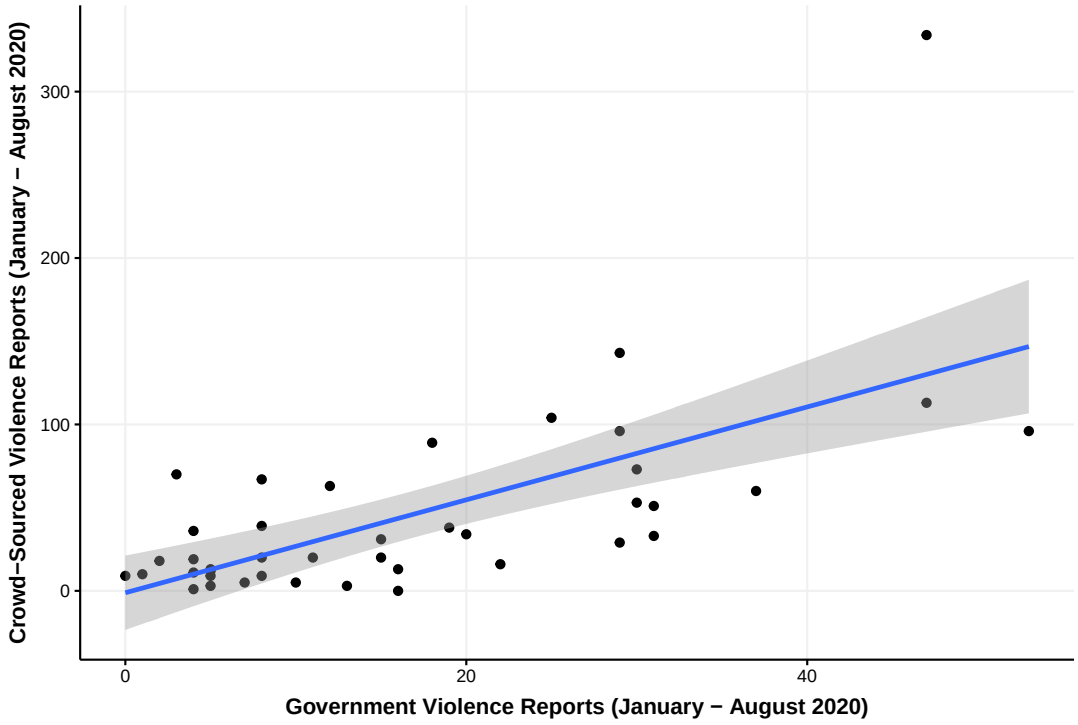

Figure 4.5: Relationship between Information from Government and Social Media Sources. - Figure plots the bivariate correlation between the neighborhood crime as reported by the social media source, Fogo Cruzado, and the government source, ISP.

Table 4.1: Covariate Balance Table

Covariate	Government Information Mean	Social Media Information Mean	P-Value
𝟙(Male)	0.4368	0.4332	0.8973
Age	4.3040	4.1771	0.0782*
Education	2.3392	2.3188	0.5378
Income	1.9536	1.9163	0.456
Race	0.4928	0.5024	0.9162
𝟙(Victimization)	0.4812	0.4959	0.7345
𝟙(Violence is Down)	0.6256	0.6457	0.4607
𝟙(Facebook Group Member)	0.5520	0.5684	0.5593
Fear of Crossfire	3.1008	3.1433	0.4417
Uncertainty Score	0.5010	0.5056	0.5687
Note:			*p<0.1; **p<0.05; ***p<0.01

Table 4.1 provides means or each of the two information treatments for relevant covariates as well as p-values to measure whether this difference is significant across the experimental conditions.

measure of uncertainty, plotted below in **Figure 4.6**.In this figure, the distribution of uncertainty scores appears to be fairly consistent across treatment conditions, with uncertainty among respondents in both the government information and social media information groups concentrated between .4 and .6. This means that most respondents reside somewhere in the middle between highly uncertain (1) and very certain (0) about the incidence of violence in their neighborhoods. There are some minor differences between the two groups (for example, those in the social media condition are slightly more likely to be uncertain), but these differences are insignificant. In addition, based on the t-tests, respondents who received the government information treatment are slightly older than those who received the social media treatment. As a result, I controlled for age in all models run throughout this chapter.

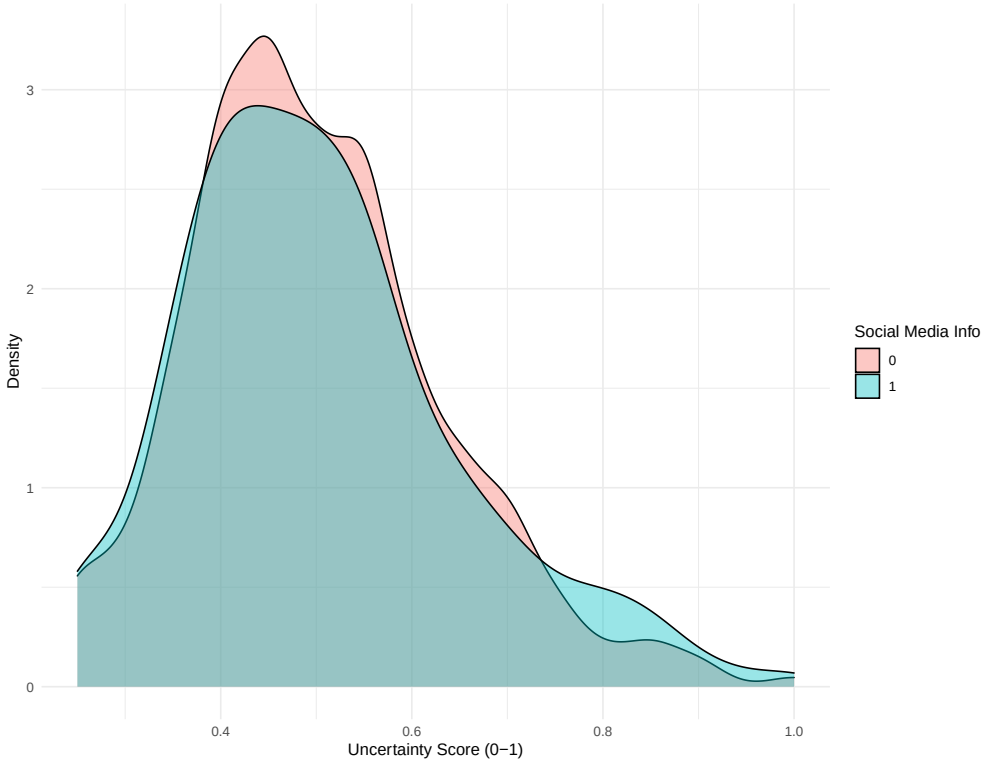

Figure 4.6: Uncertainty Scores by Treatment Assignment. - **Figure 4.6** details the density of uncertainty scores by treatment assignment. I observe no imbalance on this variable, which is critical for my analysis.

Dependent Variables: Anxiety about Violence and Support for the Police

At the beginning of the survey, prior to receiving the information from either the government source or the social media source, I asked respondents to answer the following questions:

> [**Outcome 1 (Anxiety)**] – How worried would you be about urban crime in your neighborhood, ranging from 0-100, with 0 being "not at all worried" and 100 being "extremely worried."
>
> [**Outcome 2 (Political Attitudes)**] – How likely are you to support increased police presence in your neighborhood, ranging from 0-100, with 0 being "extremely unsupportive" and 100 being "extremely supportive."

Outcomes range on a slider scale from zero to one hundred, with zero being low and one hundred being high. Respondents are required to slide a bar to record their preference. These questions provided me with a "baseline" measure for anxiety and political attitudes, my two dependent variables.

Following the treatment assignment and information prompt, I then asked these questions to respondents again (an "endline" measure). I used the endline responses for both questions as my dependent variables for this study. Following insights from Allison (1990), I control for the baseline responses as a coefficient in my regression models.[1] This measure allows me to observe the difference in the effect of new information from either the government or social media source on beliefs, in a way that is not explained by baseline perceptions. The equation to calculate the average treatment effect is:

[1]For a detailed discussion, including a simulation, of when it is appropriate to control for pre-treatment outcomes vs. calculate a change score between endline and baseline outcomes, see the `DeclareDesign` blogpost, "Use change scores or control for pre-treatment outcomes? Depends on the true data generating process."

$$\begin{aligned}\text{Political\_Attitudes}_{post(i)} =& \alpha + \beta_1 \text{Treat}_i + \beta_2 \text{Age}_i + \beta_3 \text{Political\_Attitudes}_{pre(i)} + \\ & \beta_4 \text{Neighborhood\_Violence}_i + \epsilon\end{aligned} \tag{4.1}$$

The left hand side of the equation represents the outcome variable (in this case Political Attitudes) after participants receive information from one of the two sources. The right hand side of the equation represents the treatment status (government or social media information), with controls for age, the numeric value of the information they saw about violence in their neighborhood, and pre-treatment political attitudes measured at the beginning of the survey.

4.3 Results

The results from the information experiment are detailed in **Table 4.2**. I find no significant effect of the information type on either anxiety about community violence (model (1)) or support for increased policing (model (2)) alone. However, interacting the treatment variable with a measure of prior uncertainty (as described in **Section 4.2**) reveals a complex relationship between the main variables of interest. These findings are presented in model (2), with respect to the anxiety outcome, and model (4), with respect to the policing outcome. All models control for age (a variable that was slightly imbalanced across treatment and control conditions), the numeric value of the information respondents saw, and baseline anxiety/police support, which differ based on the dependent variable.

Specifically, I find that the null effects in models (1) and (3) are likely driven by movement in opposite directions for low-and high-uncertainty respondents who receive the social media information. High-uncertainty respondents who receive information about violence in their communities from social media show a significant decrease in both their anxiety and support for increased policing. A one unit change in uncertainty for those who receive information

Table 4.2: Information Experiment Results

	Dependent variable:			
	Anxiety		Police Support	
	(1)	(2)	(3)	(4)
$\mathbb{1}$(Social Media)	−0.253	6.789*	−0.240	6.868*
	(0.993)	(3.487)	(1.182)	(4.151)
High Uncertainty		4.951		8.192
		(4.908)		(5.835)
$\mathbb{1}$(Social Media)×High Uncertainty		−13.904**		−14.081*
		(6.617)		(7.877)
$\overline{y_i}$	30.843	29.483	22.563	19.342
Observations	1,245	1,245	1,245	1,245
R^2	0.545	0.547	0.490	0.492
Adjusted R^2	0.540	0.541	0.485	0.486

Table 4.2 shows the results from the information experiment. Models (1)-(2) use "Anxiety" as the dependent variable; models (3)-(4) use "Police Support" as the dependent variable. Models (2) and (4) interact the experimental treatment variable with a continuous variable that measures respondent-level uncertainty. All models control for age, the value of the information respondents saw, whether it was higher or lower than they expected, and baseline anxiety/police support, depending on the dependent variable. Significance: *p<0.1; **p<0.05; ***p<0.01.

about violence translates into a nearly fourteen-point decrease in concerns about violence and support for increased police presence in their communities, on a 100 point scale. Low uncertainty respondents who receive information about violence in their communities from social media appear to significantly *increase* both their anxiety and support for increased policing, though the magnitude is about half the size of the uncertain group. In this case, a one-unit change in uncertainty for those who receive information about violence translates into a nearly seven-point increase in concern about violence and support for increased police presence in their communities, on a 100 point scale. By contrast, there is no significant change among respondents who receive the government information, for either high-or low-uncertainty individuals.

Disentangling these competing effects is beyond the scope of this survey experiment,

but one post hoc rationale is that in the face of new information, social media may be viewed as more trustworthy for uncertain individuals because the content comes from people "like them." At the same time, it may also reinforce existing beliefs among highly certain individuals, who can easily filter through massive amounts of information in a way that reaffirms prior beliefs. If an individual has strong prior beliefs about the given subject, they may double down on those attitudes, seek out only information that confirms their preexisting attitudes, and dismiss any other disconfirming information. Exploring these dynamics will be critical for future research on this topic.

4.4 Mechanisms

By definition, social media is distinct from other sources of information provision (e.g., television, radio, or newspaper/the periodical press) in that it has low barriers to entry, relies on user-generated content, often comes with accompanying visuals, among other relevant attributes (Zhuravskaya, Petrova and Enikolopov, 2020; Zeitzoff, 2017). Any one of these attributes might have appealled to the respondents in my survey who received information about violence from a social media platform. In the real world, information provided through this type of medium is by definition a bundled treatment. For social media information, there is a limited delay between the occurrence of an incident and information about the incident being shared online. In addition, online information about the incident may be more likely to come with photos or video, due to the fact that someone directly involved in or near the incident is the one sharing the information. Finally, the information is more likely to come from a non-expert source, in light of the information-generating process. By contrast, other types of information that comes from more official channels may have time delays for processing, could be generated after the fact, and originate from an "official" source – most often, the government. This reality is reflected in the nature of the information prompts shared with respondents in the information experiment detailed above.

The crowd-sourcing prompt used in my survey emphasized that information came from residents and was reported in real time. The government information prompt emphasized that the information came from the civil police and was posted a month later. The final variation in the prompt (gunshots vs. homicides) was required in order to ensure that I shared **accurate** information, but if anything, it would dampen the significance of the social media prompt due to nature of the violence being discussed (gunshots, as opposed to homicides). Based on the information experiment, the community-based information appears to have resonated with survey respondents more that the government provided information, leading them to reevaluate not only their political attitudes but also their perceptions of insecurity. Still, the nature of this bundled treatment makes it difficult to distinguish which aspects of the treatment are driving this change.

In order to further explore the underlying mechanisms behind the information experiment findings, I ran a separate conjoint experiment in which I showed respondents a mock Facebook post with information either about violence or about public services. The goal here is twofold: (1) to understand which *features* of social media information may lead respondents to trust one set of information over another; and (2) to determine whether any effects are specifically about violence–a highly salient topic in Rio de Janeiro–or if they might extend as well to more benign topics such as the provision of government services.

To better understand why social media drives individuals to view it as more trustworthy, I use a conjoint experiment with four attributes ranging from two to four levels to evaluate the role of trust. I showed respondents eight different pairings of these attributes, and they were required to choose which of the two pairs they would be more likely to trust based on these randomized attributes.[1] Importantly, with respect to the second question about the information type, I showed respondents social media messages that were both about violence and about the provision of public services (**Figure 4.7**). In rounds 1, 3, 5, and 7, respondents saw information about local development, and in rounds 2, 4, 6, and 8 they saw

[1] As a feature of the conjoint experiment, I required that at least two of the four attributes be different from eachother.

Figure 4.7: Conjoint Experiment Prompt

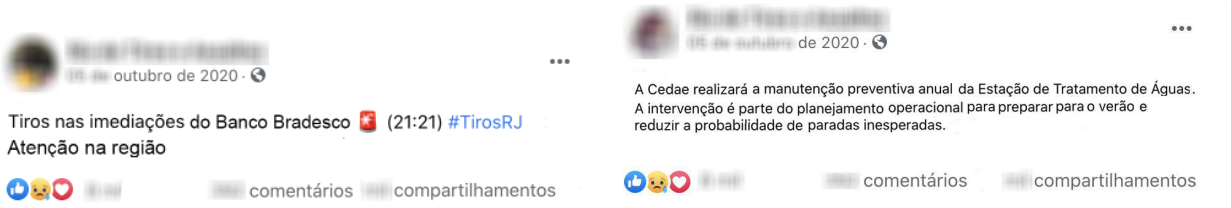

(a) Violence Prompt **(b)** Government Services Prompt

Figure 4.7 provides an overview of the two social media prompts used in the conjoint. Translated the violence prompt (a) says: "Gunshots near the Bradesco Bank (21:21) #GunshotsRJ. Attention in the area"; whereas, the government services prompt (b) says: "Cedae will perform annual preventative maintenance on the Water Treatment Depot. The intervention is a part of the operational planning to prepare for the summer and reduce the probability of unexpected stoppages."

information about local violence. The attributes and levels accompanying these images are detailed in **Table 4.3**.

Table 4.3: Conjoint Experiment

Attribute	Levels
Timing of Post	(1) *At the time of the incident/announcement*; (2) A day later
Source	(1) The person who posted: *was a member of your community*; (2) was a city councillor; (3) was a member of the media; (4) *was someone you know personally*
Engagement	(1) *The post had more than 500 reactions, shares, and comments*; (2) The post had around 100 reactions, shares, and comments; (3) The post had less than 20 reactions, shares, and comments
Media	(1) *The post included a photo of the incident*; (2) *The post included a video of the incident*; (3) The post included neither video nor photo of the incident

Figure 4.3 describes the different attributes and levels that included in the conjoint experiment. Respondents saw four different pairs for each topic detailed in **Figure 4.7** (violence or government services), for a total of eight pairs. Italicized levels are characteristics of community-based social media, which feature prominently in the bundled information treatment.

Specifically, I vary the attributes of the information by: (1) source, ranging from experts to community members; (2) speed, meaning the time between events and information de-

livery; (3) post engagement in the form of likes, comments, etc., evaluated here as a signal of community-driven support for the message; and (4) evidence associated with additional information provided by videos or photos. I have italicized the levels that emphasize attributes characteristic of community-based social media (including those which vary in the information experiment detailed above). For this conjoint, I require that each pair has at least two attributes with different levels assigned to them in the randomization. This ensures that (1) participants do not have to choose identical between pairs; and (2) there is enough variation between the pairs for participants to make a clear choice. The purpose of this design is to assess what about social media seems to resonate with citizens when they are evaluating which information to trust and to determine whether or not these findings might generalize to less salient topics beyond violence.

To analyze results from the conjoint experiment, I follow Hainmueller, Hopkins and Yamamoto (2014) in estimating the average marginal component effects (AMCEs), which calculates the average effect of each attribute in the profile on trust (my outcome), while holding all of the remaining components of the profile constant. The equation for this calculation is detailed below:

$$\begin{aligned}\text{Trust}_{ijk} =& \beta_0 + \beta_1[\text{Day Of}_{ijk}] + \beta_2[\text{Source}_{ijk} = \text{Government}] + \beta_3[\text{Source}_{ijk} = \text{Media}] \\ &+ \beta_4[\text{Source}_{ijk} = \text{Friend}] + \beta_5[\text{Likes}_{ijk} = 500] + \beta_6[\text{Likes}_{ijk} = 100] \\ &+ \beta_7[\text{Photos}_{ijk}] + \beta_8[\text{Video}_{ijk}] + \epsilon_{ijk}\end{aligned} \tag{4.2}$$

In this equation, each attribute level acts as a coefficient in the model; with *i*, *j*, and k representing the individual, choice, and alternative choice, respectively. In this design, I make several assumptions. I assume that: (1) the respondents do not use information from an earlier choice task to color their assessment; (2) the order of seeing the profiles does not matter; and (3) profiles are sufficiently randomized.

In order to guard against any potential effects of the information experiment on the

conjoint task, I ran the conjoint experiment first in the survey, immediately following the baseline demographic questions. In total, the models have 19,920 observations, representing 1,245 people who compared two different profiles eight times ($1,245 \times 2 \times 8 = 19,920$). Baseline levels for each attribute are as follows: "Less than 20 reactions, shares, and comments" (engagement); "Neither photo, nor video" (media); "A member of the press" (source); and "A day later" (timing of post).

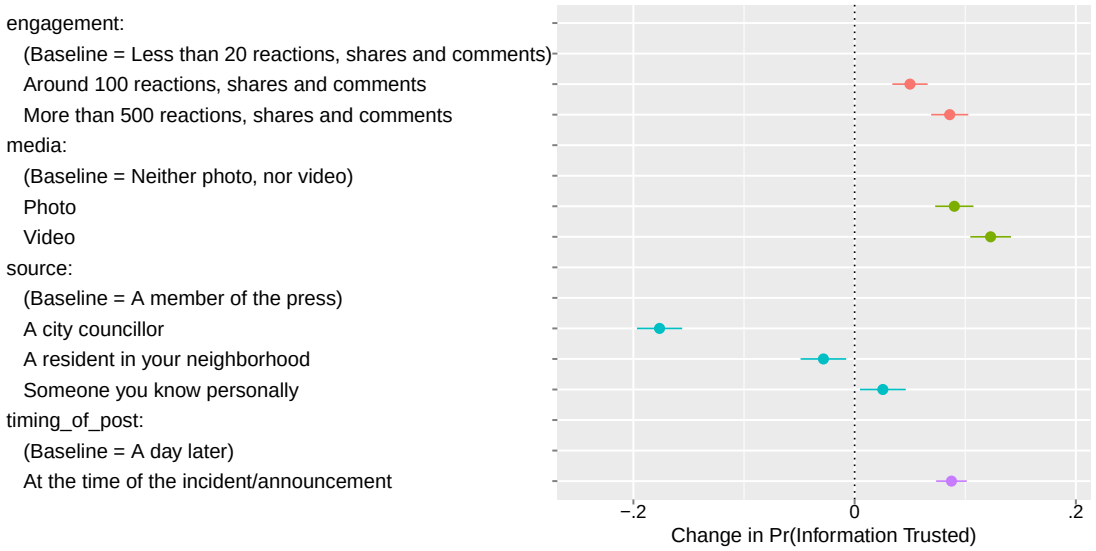

Figure 4.8: Average Marginal Component Effect - Figure plots the average marginal component effect for all attributes and levels in the conjoint experiment. Estimates that appear on the right side of the vertical line are more important for respondents when choosing between two information profiles. Estimates that appear on the left side of the vertical line are less important for choosing which profile to trust.

I plot the AMCE estimate for each attribute below, in **Figure 4.8**, clustering standard errors at the respondent level in order to account for the fact that each participant appears in the data frame eight times. As is evident from these results, a number of features tied to social media posts played a significant role in whether or not respondents decided to select a specific attribute profile. In particular, respondents trusted information more if it was provided at the time of the incident/announcement, provided by someone the respondent knew personally or by the press, had a high level of user engagement, and included external media such as videos or photos. By contrast, relative to the baseline, respondents were 17.6

percentage points less likely to trust information if it came from a government source such as a city councillor. A summary of these results is provided in **Table 4.4**.

Table 4.4: Conjoint Results Summary

	Average Marginal Component Effects (AMCE):				
Attribute	**Level**	**Estimate**	**Std. Err**	**z value**	**Pr($>|z|$)**
engagement	$\sim$ 100 reactions, shares & comments	0.050***	0.008	6.129	0.00
engagement	$>$ 500 reactions, shares & comments	0.086***	0.009	10.034	0.00
media	Photo	0.090***	0.009	10.204	0.00
media	Video	0.123***	0.009	13.117	0.00
source	A city councillor	-0.176***	0.010	-16.895	0.00
source	A resident in your neighborhood	-0.028**	0.011	-2.677	.007
source	Someone you know personally	0.026*	0.011	2.416	.016
timing_of_post	At the time of incident/announcement	0.087***	0.007	12.370	0.00

Table 4.2 shows the coefficient results from the conjoint experiment, plotted in **Figure 4.8**. Baseline levels for each attribute are as follows: "Less than 20 reactions, shares, and comments" (engagement); "Neither photo, nor video" (media); "A member of the press" (source); and "A day later" (timing of post). Significance: *p<0.1; **p<0.05; ***p<0.01.

Building on these results, I find that the importance of these attributes was surprisingly stable, regardless of whether the information was about violence or about local government services, as demonstrated in **Figure 4.9**. There are slight shifts, however, that do register across the board. For example, respondents care slightly less about timing and user engagement for posts about government services, though they still weigh them positively when determining which profile to trust. In addition, journalists are considered the most trustworthy source for news about government services, while neighbors and personal connections are viewed as more trustworthy for posts about violence. Interestingly, the single greatest factor shaping the rejection of a profile in both cases occurs when the source is a city councillor. This demonstrates that even when the government might be considered a first-hand source of information, citizens across the board have high levels of mistrust in elected officials, which may make community-based or non-government information more appealing. These results complement the findings in **Chapter 2** related to the use of social media for information about politics and trust in government across Latin America.

Finally, I evaluate this relationship conditional on levels of uncertainty, a variable central

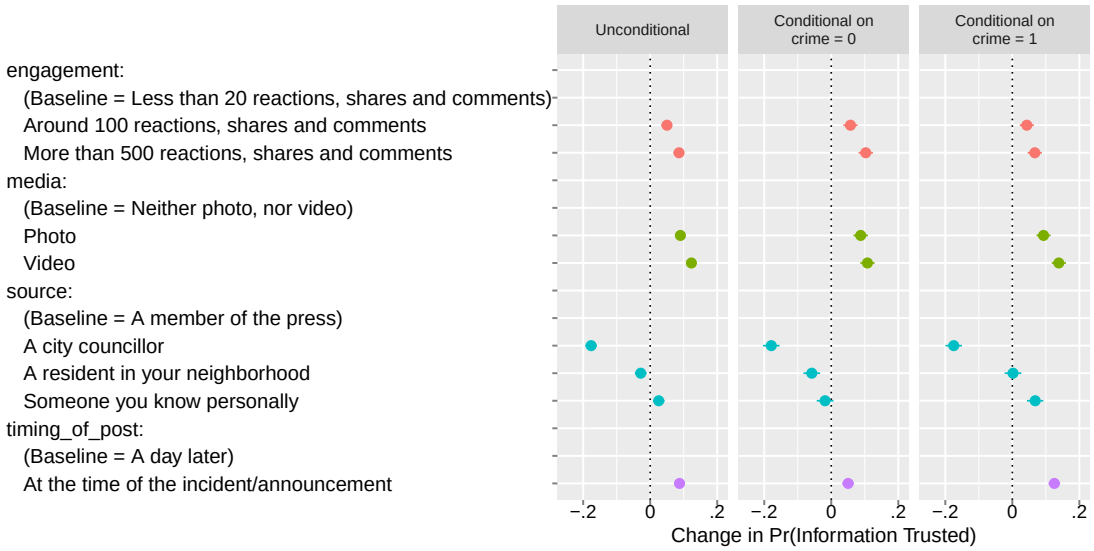

Figure 4.9: Average Marginal Component Effect by Information Type - Figure plots the average marginal component effect for all attributes and levels in the conjoint experiment conditional on the type of information – government services or violence. Estimates that appear on the right side of the vertical line are more important for respondents when choosing between two information profiles. Estimates that appear on the left side of the vertical line are less important for choosing which profile to trust. The first panel is the same as **Figure 4.8**. The second panel, "Conditional on crime = 0", show results for the government services post. The third panel, "Conditional on crime = 1", shows results for the violence post.

to my theory. In particular, I am interested in whether or not the decision to weight certain attributes as more or less trustworthy varies according to the certainty level of respondents about their knowledge of crime in their communities. For example, people who are more secure in their beliefs may be less likely to rely on user engagement as a heuristic for trustworthiness or may be more likely to view elite opinions as credible. In this context, I create a binary variable for respondents based on the continuous variable described in detail above in **Section 4.2**. Respondents with an uncertainty score of $\geq .5$ receive a 1; respondents with an uncertainty score of $< .5$ receive a 0. In total, 689 respondents are classified as "0" and 806 as "1" based on their self-reported certainty about five knowledge questions. I also restrict my conjoint sample to include *only* the paired profiles about violence. This is due to the fact that uncertainty related to violence has no theoretical link to trust in information about government services.

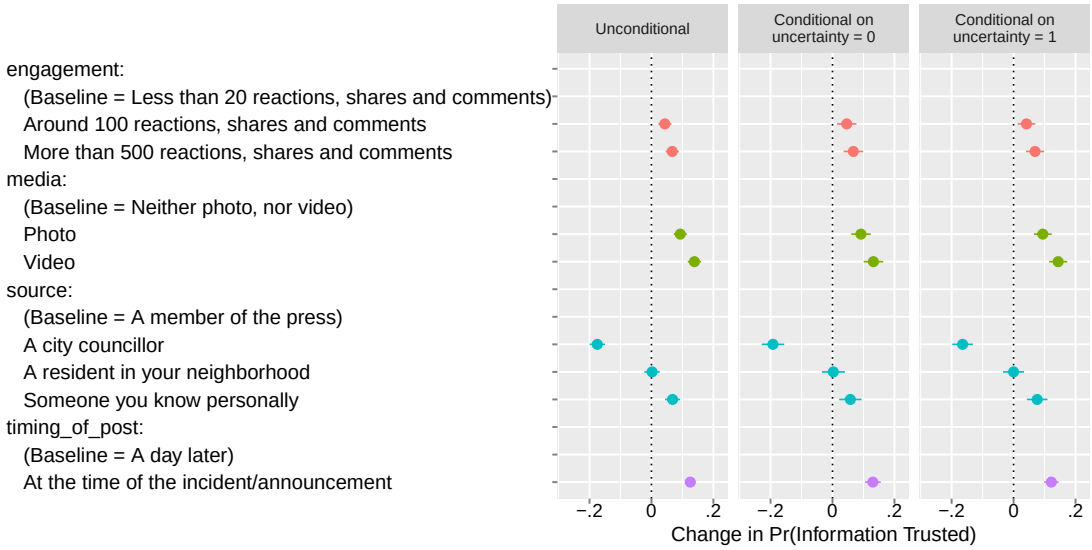

Figure 4.10: Average Marginal Component Effect by Uncertainty Levels - Figure plots the average marginal component effect for all attributes and levels in the conjoint experiment for violence-focused prompts, conditional on uncertainty – operationalized here as either "high" ($\geq .5$) or "low" ($< .5$). Estimates that appear on the right side of the vertical line are more important for respondents when choosing between two information profiles. Estimates that appear on the left side of the vertical line are less important for choosing which profile to trust. The first panel is the same as **Figure 4.8**. The second panel, "Conditional on uncertainty = 0", show results for the low uncertainty group. The third panel, "Conditional on uncertainty = 1", shows results for the high uncertainty group.

Contrary to expectations, I find that knowledge uncertainty about violence plays little role in shaping whether or not respondents weight certain attributes as more or less important when determining which profile about violence to trust. Instead, results across low-and high-uncertainty groups appear almost identical. This suggests that across the board, attributes more closely tied to social media may be considered more valuable for respondents, regardless of what they think they know about violence-related topics in their communities.

To further examine the validity of these results, I also explore how within-respondents preferences change across multiple rounds. In **Figures 4.11**, I plot the marginal mean for each round, by information type. I find some fluctuation round by round, but none of these shifts would substantively alter the interpretation of the conjoint results. Where preference differences register, they are most likely due to the nature of the information – whether it was about violence or public services – that accompanied the randomized profiles presented in the paired choice.

4.5 Conclusion

Over the past decade and a half, the information environment has shifted dramatically. Where once an authoritative relationship between broadcaster and audience defined one-way information flows, a two way conversation has now emerged. Today, as Shirky (2008) notes, "The future presented by the internet is the mass amateurization of publishing and a switch from 'Why publish this?' to 'Why not?'" It is a future, he writes, where "everyone is a media outlet." This 'mass amateurization' of the information environment has shifted gatekeepers who controlled access to information and decided what was "newsworthy" from media executives, journalists and a professionalized class to anyone with a social media account. Where once reproduction and distribution costs saddled the ability of media outlets to produce large volumes of uncurated content, today that content is seemingly unlimited, and only restricted by whether or not someone happens to record it on their phone and share

Figure 4.11: Marginal Mean by Round

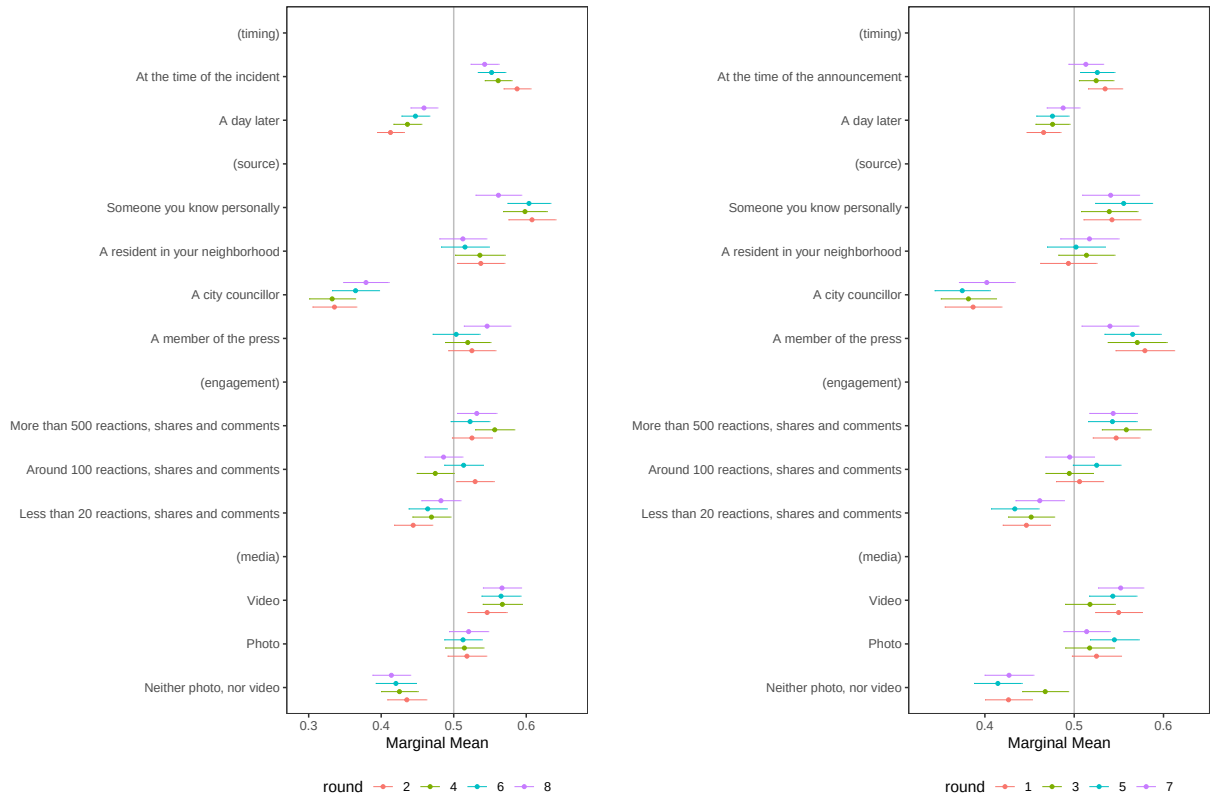

(a) Rounds about Violence **(b) Rounds about Public Services**

Figure plots the average marginal component effect for all attributes and levels in the conjoint experiment for each round. Participants saw a total of eight paired choices, four of which were accompanied by information about violence (rounds 2, 4, 6, and 8, **Figure 4.11 (a)**) and four of which were accompanied by information about public services (rounds 1, 3, 5, and 7, **Figure 4.11 (b)**).

it online.

What are the implications of this shift for public opinion formation? How does it complete theories developed prior to the incorporation of digital media into the information ecosystem? And how might these changes shape individuals' attitudes and policy preferences? Uniting prior research in this space with research in communications, I argue that two conditions will shape the effect of information on individuals: (1) how open and receptive that person is to information about the topic; and (2) whether or not it comes from a source they trust. I argue that the first condition is universal across contexts – all people will dismiss information that does not confirm their prior beliefs if those beliefs are strongly held. The

second condition, however, may vary by context, and is largely shaped by both the political and media environment of a country or community. In Brazil, where trust in government is low, and the news media is less established, information from social media may be viewed as more trustworthy because, unlike traditional gatekeepers of information, it comes from "people like them." As a result, I expect both the strength of prior beliefs and the nature of the information source – whether it comes from the community or not – to play a significant role in shaping when information leads to updating.

I test the importance of these conditions in shaping public opinion formation and attitude change through an original survey. I operationalize the openness to new information through an index that measures the strength of prior beliefs. To construct this measure, I ask respondents five crime-related questions. I follow up each knowledge question by asking respondents to rate their certainty about their answer to the prior question. I incorporate these ratings into a measure of "knowledge certainty" and hypothesize that individuals who claim higher knowledge certainty will be less likely to update their beliefs than individuals with less knowledge certainty, regardless of the information source. Information source does, however, play a critical role.

Building on findings from **Chapter 2**, I hypothesize that uncertain individuals who receive information from social media sources will be even more likely to update their beliefs than uncertain individuals who receive their information from more traditional sources such as the government. To test these effects, I randomly assign respondents to receive truthful information about violence from the government or from social media. After showing respondents information about violence in their neighborhoods, I ask them to rate their level of concern about violence in their communities (outcome 1) and their support for increased policing (outcome 2).

I find that information from social media differentially affects individuals based on their prior perceptions. In the case of this experiment, social media information leads individuals with weak priors about violence and crime in their communities to become less anxious

about violence in their neighborhood and less likely to support increased policing. However, individuals with strong priors about violence who received information from social media are more likely both to be concerned about violence and support an increased police presence in their communities. I find no such heterogeneous effects on public opinion formation for high- and low-certainty respondents who receive information about violence in their communities from a government source.

To fully understand the dynamics driving the difference in high-and low-certainty individuals who received information from social media sources is beyond the scope of this dissertation. Nevertheless, I am able to explore why individuals may be more willing to incorporate information from social media into their attitudes and policy positions. Through a conjoint experiment, I find that respondents in my sample are significantly more likely to trust information if its dissemination resembles that of social media content. Respondents are more likely to trust information that has high user engagement, us shared in real time, does *not* come from an official government source, and incorporates additional evidence such as videos or photos. These findings are not conditional on uncertainty levels and remain largely unaffected by the focus of the content. Noticeably, the single largest predictor of trust is that the content did not come from the government. This demonstrates why the crowd based nature of social media may be particularly important for public opinion formation in Brazil's second largest city. Future research should seek to disentangle the divergent findings from the information experiment and explore the dynamics of trust in social media content, which is significantly more likely to shape political attitudes in this context.

5

The Bias of the Crowd

"What you see is all there is."

– *Daniel Kahneman, Thinking Fast and Slow, pg. 85*

With the rapid expansion of Internet connectivity, citizens are now able to reach out into the world, far beyond their direct acquaintances and communities. In Brazil, social media first took the spotlight in 2013, when calls to *"Vem Pra Rua!"* (come to the streets) sparked a nationwide protest that began over an increase in bus fares.[1] Over time, the use of social media platforms has expanded and become more mainstream. Today, many urban communities utilize this type of communication medium to share stories and connect with neighbors online. Such usage of these platforms is designed not only to foster a sense of community, but also to glean local-level insights that traditional media and political elites either can not or will not provide.[2] Unlike traditional sources of information (namely, newspapers and television), in these online networks, individuals living in the community

[1]Costa, Camilla. "Brazil unrest: Protestors 'subvert' advertising slogans." *BBC News*, June 29, 2013, https://www.bbc.com/news/world-latin-america-23108688.

[2]This is a way of utilizing Facebook resembles the social media platform NextDoor, an online portal for members to share information about their neighborhoods. Unlike NextDoor, Facebook users neither have to provide evidence that they live in the community to follow the content, nor do they need to register for an account beyond their Facebook profile.

are responsible for providing, curating and promoting content. With citizens in control, how different is this type of information and what are the implications of this content for shaping political beliefs?

In the previous chapter, I showed that individuals who were initially more uncertain about the amounts of crime in their neighborhoods were more susceptible to attitude change when exposed to new information about violence. In the current chapter, I investigate the relationship between actual experience with crime and perceptions about it, which leads to some suggestions about why social media information may be systematically biased toward emphasizing threatening information.

I begin first by evaluating how mainstream media and social media talk about violence. Despite perceptions in the mainstream media that violence "sells," I show that the proportion of online conversation about violence in Facebook communities is significantly higher than the proportion of conversation about violence in mainstream news outlets. This result demonstrates that social media content is even more prone than traditional media to promoting sensationalist or anxiety-inducing content in response to humans' seemingly hard-wired tendency to give more attention to potentially dangerous things in their environments than to other types of information (Kahneman and Tversky, 1979).

Posts about violence make up 31.5% of all conversations in community pages (192,812 posts) and only 17.5% of articles in the news. I also find that information about violence is more likely to resonate in community pages, as measured by the total number of reactions to a post as a proportion of the average annual page reactions. Citizens are 28.3% more likely to react (through likes comments, shares, etc.) to posts about violence on community page posts than to other types of posts. Citizens are 19.6% less likely to react to violence posts when they are shared by the news media. This is due to the nature of social media, which prioritizes bottom-up, citizen generated content as opposed to curated, top-down information, and the nature of community pages, which are often highly localized to happenings within specific neighborhoods.

While social media pages appear to be more likely to share content about violence that traditional media platforms, I theorize that there is wide variation in the frequency with which communities share and engage with information about violence online. I argue that this variation depends on perceived exposure to threat. While we might expect citizens in communities more affected by violence to speak more frequently about it, I expect that actual levels of violence do not wholly explain where these conversations will be the most frequent or how citizens will react to them. Specifically, average people who either live in very violent neighborhoods or in very safe ones will feel more certainty about crime levels in their neighborhoods than do those who live in neighborhoods where violence levels are not at either extreme. However, in communities where citizens have less certainty about violence (whether it is typically high or low), conversations about violence will be more common and more likely to generate a response from the community. Empirically, this equates to neighborhoods in the middle-income neighborhoods, due to a high degree of correlation between neighborhood development and violence.

I test the usefulness of this suggestion by comparing Facebook conversations about violence among high, low, and middle income neighborhoods. Specifically, I assembled a dataset of 610,389 Facebook posts from 221 community pages in Rio de Janeiro shared between January 2013 and December 2019 using `CrowdTangle`, a public insights tool owned and operated by Facebook.[1] Using a combination of manual coding and natural language processing, I classified each post based on whether it is or is not about violence and calculated the total number of reactions (in the form of likes, comments, shares, etc.) to each post.I then divided these neighborhoods by socioeconomic development, as delineated by responses to questions in the national census. I show that people who live in middle-income neighborhoods post more information about the crimes that do happen than do those who are more accustomed

[1]According to Facebook, "CrowdTangle is a Facebook-owned tool that tracks interactions on public content from Facebook pages and groups, verified profiles, Instagram accounts, and subreddits. It does not include paid ads unless those ads began as organic, non-paid posts that were subsequently 'boosted' using Facebook's advertising tools. It also does not include activity on private accounts, or posts made visible only to specific groups of followers."

to crime or those who live in very safe neighborhoods.

Within the Bayesian framework for understanding attitude change, frequent posting about crime means that people in middle-income neighborhoods who participate in Facebook neighborhood pages are exposed to more messages about crime than are those who participate in neighborhood pages in more violent neighborhoods. This finding thus suggests that people who live in neighborhoods with middling amounts of violent crime come to believe that crime is more prevalent than it really is because they see many more posts about the relatively small number of crimes that actually happen. Although these results represent only a snapshot in time, this finding helps to explain how crowd sourcing can lead to systematic bias and exaggerated perceptions about threats rather than convergence on an accurate understanding of the world. Due to the higher trust attributed to content generated on social media (as demonstrated in **Chapter 4**), this may play a critical role in shaping their attitudes and beliefs, particularly about issues like violence.

In **Section 5.1**, I provide a brief summary of my theory and hypotheses for this analysis, drawing on the more detailed theory outlined in **Chapter 3**. In **Section 5.2**, I provide additional background on the process of identifying community Facebook pages, as well as details on the geographic scope and membership size of these communities. I also discuss how I identified and coded violence posts, which serve as the basis for my independent variable, the proportion of posts about violence per month by source. In **Section 5.3**, I compare the monthly proportion of and reaction to posts about violence in the mainstream media and on neighborhood community pages. Finally, in **Section 5.4** and **Section 5.5**, I evaluate the association between social development at the neighborhood level and the number of and reactions to posts about violence, respectively.

5.1 Theory Overview

More information – and in particular more trusted information – should in theory contribute to making the public as a whole more informed. This should, as Bayes assumed, bring public opinion closer to the reality that generated the information. But what if the information is inaccurate or prone to its own biases? As demonstrated in **Chapter 4**, in low government-trust settings, citizens may be more willing to trust information on social media, because the "mass amateurization" of the information environment means that content comes from people "like them," who likely share their concerns, interests and experiences.

Yet information that comes from "the crowd" may also suffer from the biases of the crowd, which will likely differ across different contexts and from person to person in less systematic ways than the known biases built into the mainstream media. With the crowd in charge of what content gets shared, the reputational risks associated with sharing biased or inaccurate content are largely eliminated. What is newsworthy in some communities may be vastly different from what is newsworthy in other communities; what is shocking in some places, may barely register in other places; and what draws a reaction in some contexts, may matter very little elsewhere.

With respect to violence, these biases may be particularly acute, and shaped by prior experience with and exposure to violent incidents. Some communities may use social media as an opportunity to counter narratives about persistent violence, while others may share this content because it is shocking, but rare. What is salient in some contexts, may not be in others. But because community members are content producers, the content that is created and disseminated may reflect the biases of the community in unexpected and reality-skewing ways. In other words, information will not necessarily lead to a more informed public, due to distortions in the information environment, as shaped by the preferences and perceptions of content creators. These distortions, which may not accurately reflect the on-the-ground reality, may haves implications for attitude formation and policy preferences due to the

greater willingness of citizens to place their trust in this content.

I explore the nature of these biases in two ways by comparing: (1) mainstream media outlets and community pages; and (2) community pages across different levels of development. The hypotheses I test in this chapter include:

H1) Social media content will be more likely to cover shocking or fear-inducing content than mainstream media.

H2) Citizens are more likely to engage with shocking or fear-inducing content from social media than mainstream media.

H3) Shocking or fear-inducing content on social media will be more common in places where citizens have higher uncertainty about the frequency of these events.

H4) Citizens will be more likely to engage with shocking or fear-inducing content on social media in places where citizens have higher uncertainty about the frequency of these events.

5.2 Data

5.2.1 Identifying Media Outlets and Facebook Community Pages

Using online data collection, I explore the association between the frequency of posts about violence (and reactions to those posts) in major media outlets and on Facebook community pages. In addition, within the community pages, I evaluate how this relationship varies based on neighborhood characteristics, such as the Social Development Index (IDS) and actual violence levels.[1] To conduct this analysis, I have collected Facebook posts from 19

[1]To construct the Social Development Index (IDS) eight indicators were used from the 2010 IBGE Demographic Census. These indicators include: (1) the percentage of permanent private households with adequate water supply; (2) the percentage of permanent private households with adequate sewage; (3) the percentage of permanent private households with garbage collected directly by a cleaning service or placed in a cleaning service bucket; (4) the average number of bathrooms per resident; (5) the percentage of illiteracy

major media outlets and 221 community pages covering 126 out of Rio's 160 delineated neighborhoods. The original list of 156 Facebook pages analyzed in Nunes (2017) provides the basis for this list, which I expand systematically.

To identify media outlets, I rely on a list of the most popular news sources in Brazil, as compiled by the Reuters Digital News Reports, which collects annual data on media habits in a select number of countries.[1] Importantly, I assume that the posts on the Facebook pages of these companies features the articles and news segments that they share on their primary platforms (e.g., the radio or television). I validate this by selecting a random sample of 1,600 posts to confirm that the link provided in the body of the post connects to a formal article written by that media company.

In a geographically representative, city-level survey of 1,245 citizens that I conducted, 56% said that they were members of a local community Facebook group. To identify these local community pages, I searched for the neighborhood name in combination with common words for denoting community pages on Facebook. These include: (1) *alerta* (alert), (2) *presente* (present), (3) *noticias* (news), (4) news, (5) *viva* (live), (6) *fala* (talk), (7) *associação de moradores* (association of residents); and (7) *depressão* (depression, a term used to ironically talk about neighborhood problems).[2] I then pulled all the pages that appear in Facebook's `CrowdTangle` search feature and that are relevant to the project. Due to heightened data restrictions, `CrowdTangle` is now the main way for researchers to access Facebook content. Through `CrowdTangle`, researchers have access to the entire universe of public Facebook pages and groups and can pull historical post information since the group

among residents aged 10 to 14 as compared to all residents aged 10 to 14 years; (6) per capita income of permanent private households; (7) the percentage of private households, with per capita household income up to one minimum wage; and (8) the percentage of private households, with per capita household income more than five times the minimum wage.

[1]New sources include: Band FM, BandNewsTV, CNNBrasil, Diário do Rio de Janeiro, Estadão, Folha de S.Paulo, G1, GloboNews, Journal da Record, Jornal Extra, Jornal O Dia, O Globo, O Povo Online, Rede Globo, RedeTV!, Rádio BandNews FM, Rádio Rio de Janeiro, SBT Terra, and UOL. A complete list of all news media outlets, followers as of December 14, 2020, number of posts, and whether they are local to Rio is included in **Table 7.2**.

[2]A complete list of all the pages, their date created, number of posts, and total followers as of December 14, 2020 is provided in **Table 7.1**.

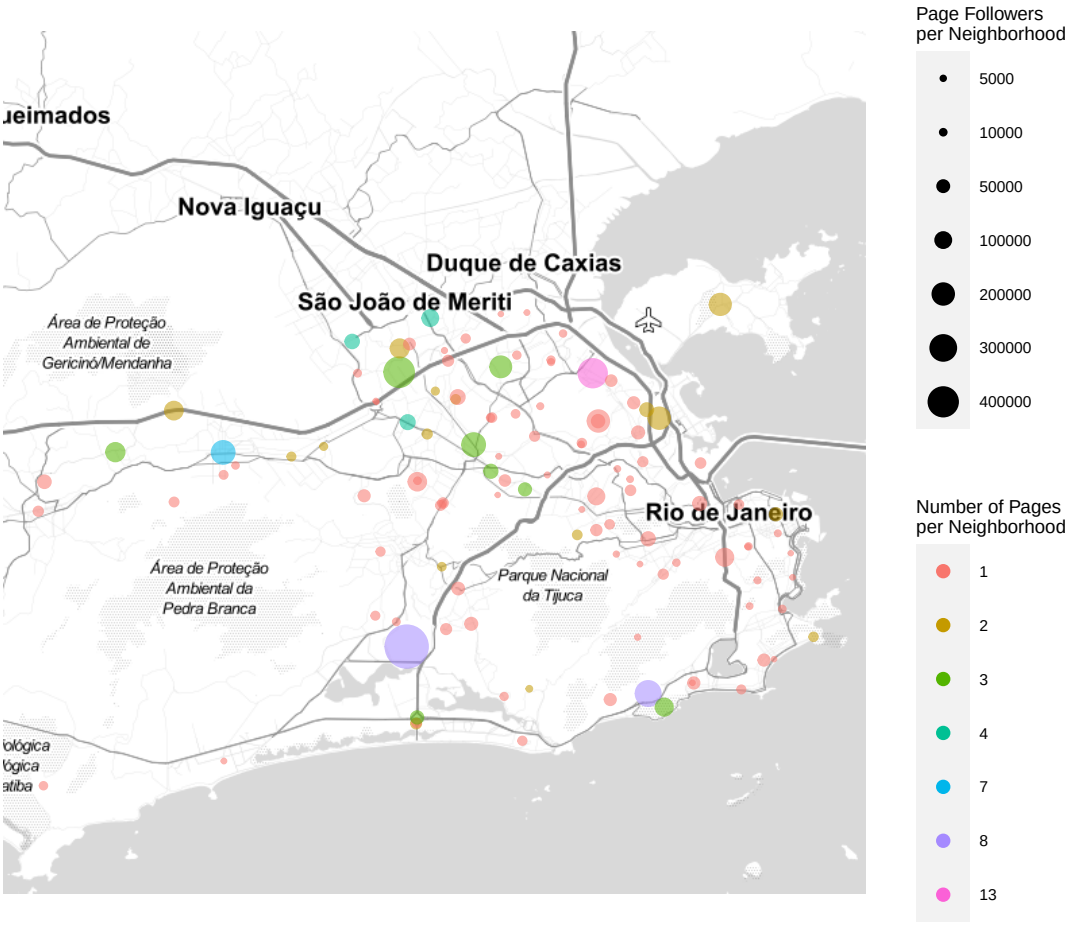

Figure 5.1: Geographic Location and Community Size of Facebook Pages - The figure shows the geographic location of Facebook community pages in my sample, the total number of groups per neighborhood, as well as the total combined membership of all pages in that neighborhood. I geocoded neighborhoods at the centroid, using the `ggmaps` package.

began, as well as information about reactions, including numbers of comments, shares, likes and other reaction types for each individual post.[1] I geocoded the pages using the centroid of the neighborhood that they cover and plotted the location of each page.

I plot the location of these pages, the number of distinct pages, and the total number of followers per neighborhood in **Figure 5.1**. Locations are geocoded to the centroid of the specific neighborhood using the `Google Maps API`. As is evident from the figure, there is significant geographic dispersion of the pages in my sample across the municipalities. Although due to data restrictions I am unable to identify the number of unique members, in total, these pages reach nearly nine million users, with an average of 40,4866 users per page. Some neighborhoods have multiple pages, including Vila da Penha in Rio's North Zone, which is a low-income community with 13 pages.

Following the identification of community pages, I filtered out pages with fewer than 2,000 followers and those that had not posted in 2020.[2] According to Reuters Digital News Report, Facebook is the second most popular social media site (behind WhatsApp) in Brazil. However, according to my geographically representative, city-level survey, it is by far the most popular source for obtaining community-level information.[3] **Figure 5.2** plots the growth of these Facebook communities over time, by year, along with key events that are commonly linked to the expansion of social media in Brazil, including the Vem Pra Rua protests in 2013, where social media was used to organize, and the World Cup and the Olympics, where social media was specifically used in the absence of government support for community safety, and

[1]In addition to examining Facebook pages, which is the focus of this chapter, I explored the content in Facebook groups, which also exist across Rio's 160 neighborhoods. Unlike pages, these public groups operate much more for the purpose of buying and selling used goods, reporting lost pets, etc., and do not, for the most part incorporate the sharing of community information or news as a main objective. In total, I evaluated the content for 47 such groups.

[2]In addition to the neighborhood-level, and to verify that I am identifying groups at the correct level of disaggregation, I conducted the same search process on the "Human Development Unit," which comprises 1,136 territorial areas that are subunits of the neighborhood level and frequently used for the census. Of the 201 pages, only four additional pages were recovered through this process. This helps to ensure both somewhat widespread influence and contemporary relevance.

[3]In the survey, I asked participants to rank six different sources from most used to least used for information about their neighborhood. Options included: 1. Television; 2. Newspaper; 3. Facebook; 4. Radio; 5. WhatsApp; 6. Twitter; and 7. Your neighbors, friends or relatives. Nearly 1 in 3 respondents selected Facebook first; followed by WhatsApp, television and neighbors/friends/relatives, respectively.

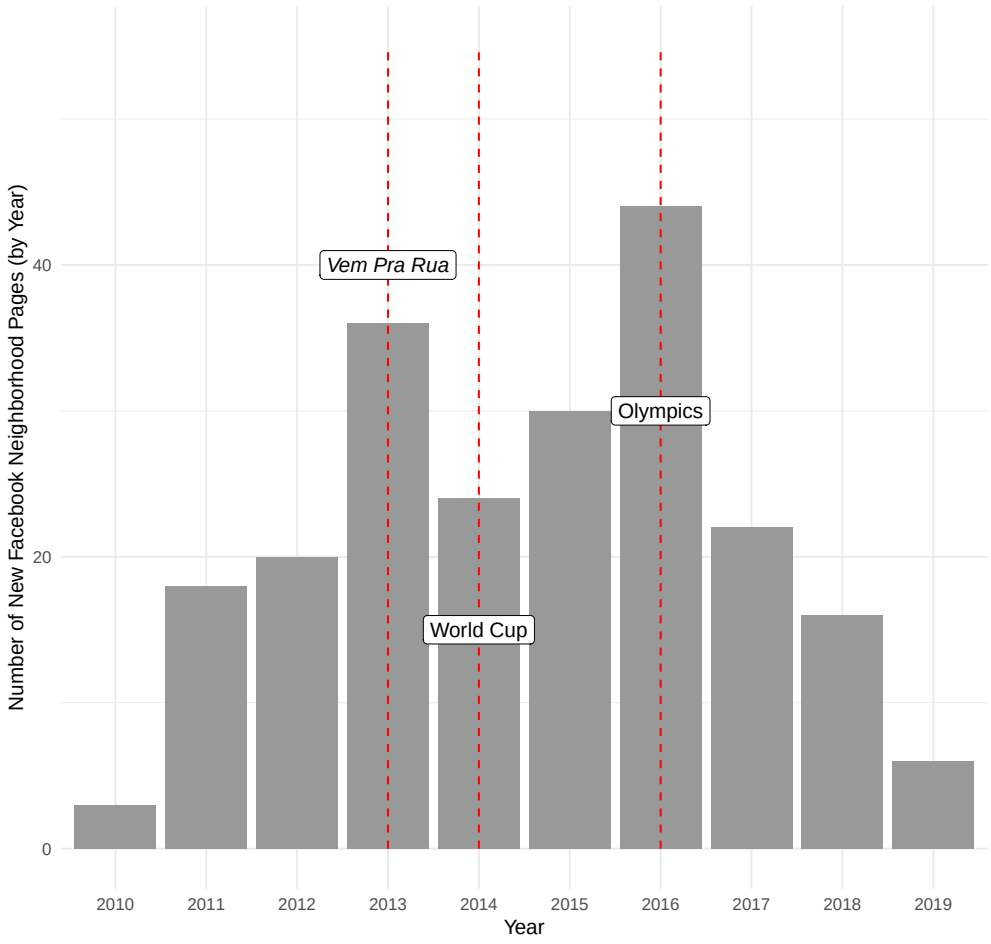

Figure 5.2: Facebook Pages per Year - This figure shows expansion of Facebook pages over time, by year. It also includes events that are considered popular catalysts for the expansion of social media in Brazil, including the *Vem Pra Rua* protests in 2013, the World Cup in 2014, and the Olympics on 2016.

was instead largely focused on the optics of the mega-sporting events.[1]

5.2.2 Coding Violence Posts

In total, I coded the content of nearly two million posts extracted from public Facebook pages using `CrowdTangle`. These posts straddle two different types of media: (1) established media channels, where content is curated by media executives and may cover local topics on occasion, but is more likely to cover broad trends; and (2) crowdsourced community pages, where content is curated by the followers who post and is more likely to be localized.

Using text analysis, I categorized posts as "violence-focused" using a dummy variable if the text of the post includes mention of topics related to gunshots, homicides, police operations, assaults, gang-related topics, robberies, militias, kidnappings, or just violence in general. I identified these topics from a list of crime-related statistics compiled by the Institute for Public Safety (ISP), a state-run institution in Rio whose purpose is to document all reports of violence throughout the state. To classify posts, I proceeded in two steps. First, I identified words clearly associated with violence, as well as all of their synonyms (for example, for assault, I looked for the words "assalto", "ataque", "estupro", and "extorsão").

As a second step, I used Natural Language Processing to identify other words in my Facebook post corpus that are used in ways similar to the more obvious crime-related words. Using the `word2vec` algorithm, which utilizes deep learning through neural networks on a large text corpus to identify word synonyms, I identified other related words not previously captured in the key word approach.[2] In total, I used 106 different key words compiled through these two processes to classify posts as either "about violence" or "not about violence." To evaluate the accuracy of this procedure, I randomly sampled 1,600 posts and found 58 posts from my sample that were misclassified.[3] For a full list of search terms, see **Appendix**

[1] For more on the legacy of mega-sport events, and in particular, the Rio Olympics, see: Talbot (2019).

[2] In this process, I incorporate all related words in the top 50 matches based on cosine similarity for the terms "crime", "violencia", and "homicidio" that were not previously included in the key terms list. More details on the `word2vec` algorithm can be found here: https://code.google.com/p/word2vec22.

[3] Often, this misclassification occurred when, for example, a post was about a celebrity who had died,

Table 7.3.

I evaluated these data at two different levels depending on the analysis. To evaluate the relationship between mainstream media and community pages with respect to posts about violence and reactions to these posts, I used the Facebook data at the level of the post, where each observation is classified based on whether or not it pertains to violence and whether it comes from an official news source or a community page. To evaluate the association between posts about violence in community pages and neighborhood attributes, I aggregated this data to the CISP-level for every month-year.[1] This also allowed for comparisons with official crime statistics from the city. In this analysis, I calculated the proportion of posts about violence, as compared to all posts, in the CISP-month-year as my main dependent variable. In using the proportion, I ensured that any effects I identified were not due to the fact that some groups post more often than others, but rather due to the amount of attention violence received as a proportion of all conversations in a community. I calculated the proportion of reactions to posts about violence, as compared to all reactions, at the CISP-month-year level for this analysis level as well.

5.2.3 Neighborhood Level Data

Finally, I relied on several different sources to collect data at the neighborhood level. First, I collected data for ISP's online data portal, which tracks violent crime incidents at the CISP level on a monthly basis. The CISP level is connected to the civil precinct, and there are 39 CISPs located in the municipality of Rio de Janeiro. Each CISP covers on average approximately 162,000 people and four neighborhoods.[2] The CISP level represents the lowest publicly-available, geographic-level crime indicators. To compile data on neighborhood level attributes, I utilized census level data to calculate the average age, development index (known

but not through urban violence.

[1]The CISP level is the lowest available level for which I was able to find both violence statistics and socioeconomic statistics. Frequently, each CISP encompasses several neighboring communities.

[2]The ISP Database is available here: https://www.ispdados.rj.gov.br:4432/.

as the Social Development Index), and other relevant neighborhood level characteristics of the CISP area.[1]

5.3 Violence on Community Pages and in the News Media

In this section, I assess the relationship between post origin (mainstream media vs. community page), whether or not the post is about violence, and reactions to posts about violence. Historically, the media has been criticized for covering sensationalist topics with higher frequency (Altheide, 2002), which stems from both consumer demand and the economic business model of these companies (Trussler and Soroka, 2014). As Shirky (2008) notes, this is the inherent tension within profit-seeking media outlets, which when choosing between different types of content, must decide whether they would prefer to "save souls" or "sell soap." This is particularly true in Brazil, where ratings-grabbing has been blamed for the increased anxiety and fear about violence among urban citizens (Contrera, 2002; Ramos and Paiva, 2007; Ramos, Paiva and Nunes, 2013). Given the nature of social media – the fact that it is citizen curated and can be hyper local – it is possible that these platforms further encourage fear-mongering, even compared to sources that are already criticized for their promotion of sensationalist news. I evaluate this possibility below.

In total, between the years of 2013 and 2019, I identified 192,812 posts that referenced topics related to violence out of a total of 610,389 posts. The proportion of posts in community pages about violence over time are plotted in **Figure 5.3(a)**. This represents 31.5% of conversations over the seven year period covered by this research that relate to violence or nearly double the proportion of coverage in the news media (**Figure 5.3(b)**).

[1]When data is provided at the neighborhood level (a smaller unit of analysis than the CISP), I average the values based on the neighborhoods covered by the CISP (e.g., CISP 10 covers Botafogo, Humaita, and Urca). This data is available here: https://www.data.rio/datasets/%C3%ADndice-de-desenvolvimento-social/data.

Figure 5.3: Posts About Violence per Year (2013-2019) on Community Pages and in Media Outlets

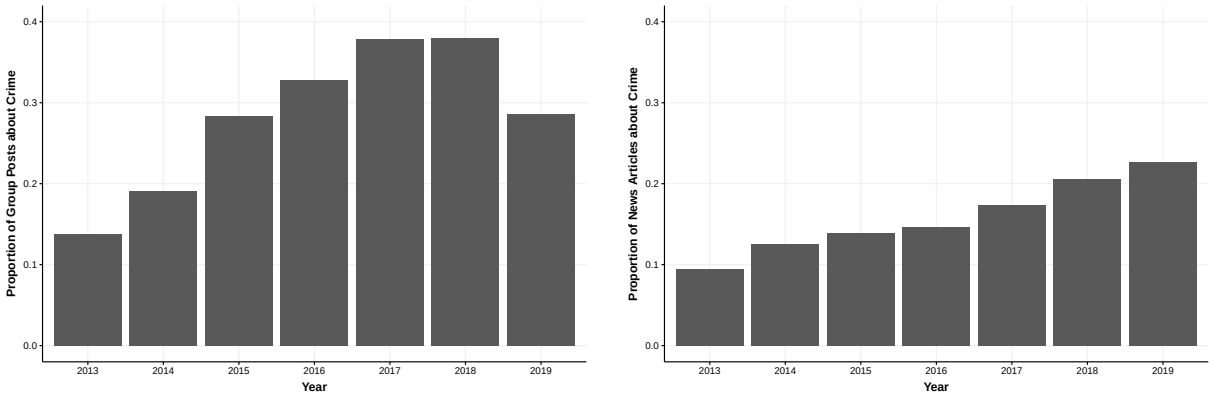

(a) Community Pages **(b)** Media

Figures shows the proportion of violence-related Facebook posts per year from neighborhood Facebook pages (a) and News Media (b) between January 2013 and December 2019. The post is coded as a "1" if it references any of the topics and a "0" otherwise. Posts are collected using `CrowdTangle`.

With respect to mainstream media between the years of 2013 and 2019, of the 1,353,989 posts from news outlets, 237,525 covered topics related to violence. During this time, approximately 17.5% of all posts on these mainstream media pages focused on violence-related topics.

In line with **H1**, I find that the proportion of posts about violence in the media is significantly lower than the proportion of posts about violence on community pages. This may be due to a variety of factors, including the curated nature of news media, the hyper-localized content of community pages, or the national orientation of news media sources.. Regardless, it is clear social media may be even more prone than traditional media to consistently peddle shocking stories – frequently, with ties to violence. Coverage of violence is less likely in mainstream media and more common on local social media. Social media may be usurping this role by providing a direct platform where citizens can share their experiences with violence and react to the experiences of others.

In line with **H2**, I find that citizens react less strongly to news media posts about violence

Figure 5.4: Reactions to Posts that Are and Are Not About Violence in Community Pages and Mainstream Media

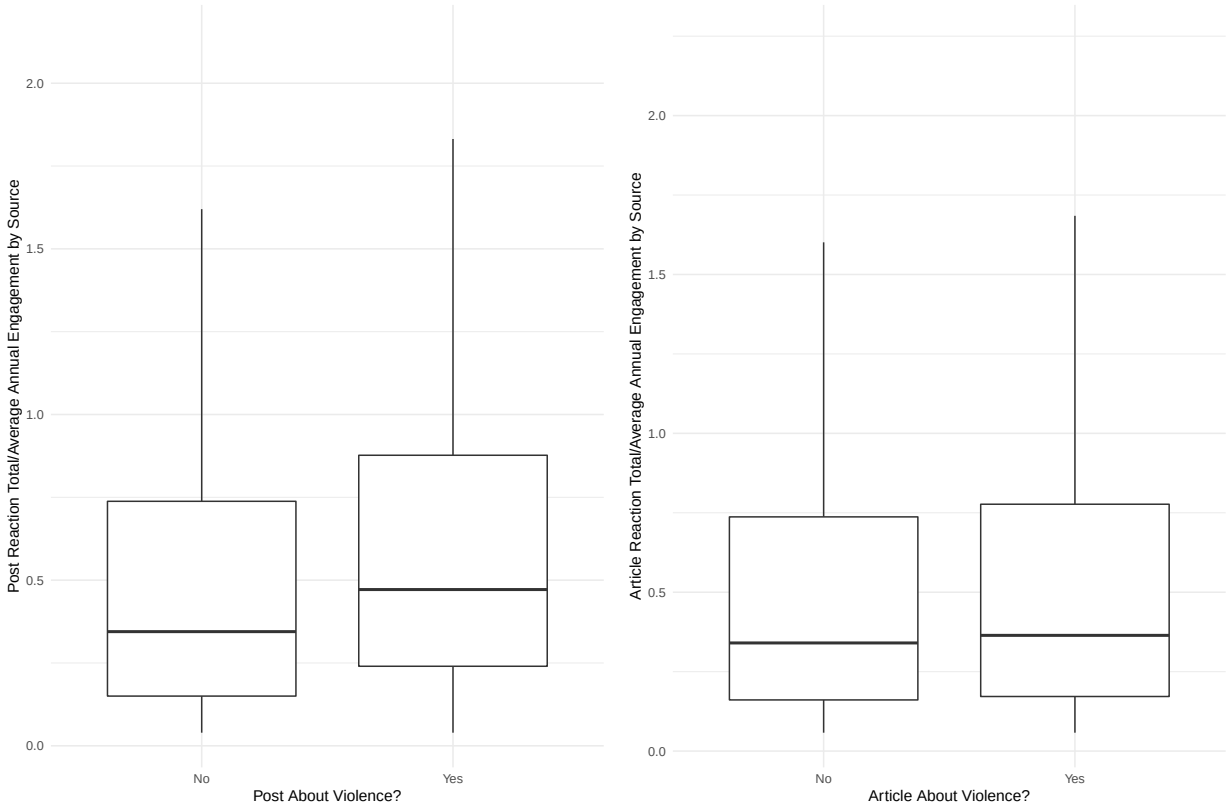

(a) Community Pages **(b)** Media

Figures show the interquartile range for average reactions for posts that are and are not about violence in community Facebook pages (a) or News Media (b).

than to community page posts about violence.[1] This is evident both in **Figure 5.4**, which plots the interquartile range for post reactions about violence and not about violence, from both community pages (a) and the media (b). In the media, articles appear to garner similar reactions regardless of subject matter. This does not appear to be the case on community pages, which are more likely to see user engagement in posts about violence. **Table 5.1** confirms this effect: in particular, the interaction effect in model (2) is negative and significant, illustrating that posts about violence in the news media are significantly less

[1]I quantify reactions as the sum of comments, likes, and shares, as well as wow, haha, anger, and sad reactions per post, divided by the average total reactions for a post from that source in a given year. As a concrete example, a single post on February 20, 2019 from *Folha de São Paulo* had 200 total reactions. The average post from that source in 2019 had 150 reactions. Therefore, the "reaction" variable is $\frac{200}{150} = 1.33$).

likely to garner reactions than posts about violence on community pages. It is possible that users compartmentalize the sources of information about violence and view information that is more local as more relevant and potentially more trustworthy or valuable, as discussed in **Chapter 4**. As a result, they engage with it at higher rates than they do with more official sources such as the news media or the government.

Table 5.1: Relationship between Post Reactions, Content and Sourcing

	Dependent variable:	
	Specific Post Reaction / Average Annual Reactions	
	(1)	(2)
$\mathbb{1}$(Violence Post)	0.165^{***}	0.283^{***}
	(0.024)	(0.025)
$\mathbb{1}$(Source = News)	0.022^{***}	0.074^{***}
	(0.003)	(0.010)
$\mathbb{1}$(Violence Post)$\times\mathbb{1}$(Source = News)		-0.196^{***}
		(0.031)
Observations	1,964,378	1,964,378
R^2	0.0004	0.001

Table 5.1 show the relationship between source (News Media vs. Community Page), reactions (total reactions as a proportion of average annual reactions by source), and post content (violence vs. not violence). Models use both month and year fixed effects to account for both season trends and annual changes. *p<0.1; **p<0.05; ***p<0.01.

5.4 Community Pages, Violence, and Social Development

Although the frequency with which community pages talk about violence is much greater than the frequency at which these topics appear in the mainstream media, it is still possible that this content is not spread evenly across the city and thus reflects trends at the community

level. Given the localized nature of community pages, as well as the important role these pages play in promoting community engagement, a logical follow-up question is whether these conversations about violence reflect actual patterns on the ground or are distorted in some way by the nature of the community where the information is being shared. In particular, how does the social development of a neighborhood affect how citizens experience and share information about violence in their communities?

I evaluate the relationship between social development and two different measures of violence – (1) the per capita number of posts about violence on community pages and (2) the per capita number of violent crimes registered by the Public Security Institute. If community pages are accurately representing the on-the-ground reality of citizens, then they should track together. However, as is evident in **Figure 5.5**, this does not seem to be the case. Social development is positively correlated with the per capita number of Facebook posts about violence and negatively correlated with the per capita number of violent crime reports filed through the Civil Police.

I find that a one unit increase in the IDS leads to a significant decline in violent crime reports per capita by neighborhood in a given month-year. This is to be expected – as development increases, a community is less likely to see violence. I also find that a one-unit increase in the IDS leads to a corresponding increase in the number of posts about violence per capita on community pages. Both of these effects are significant, though this evidence is far from conclusive. However, it does provide suggestive evidence that the posts on community pages may not be wholly reflective of the urban violence landscape. As such, this allows me to examine *why* the two deviate, *where* the proportion of Facebook posts about violence is highest, and *what* neighborhood attributes might contribute to this diversion. I explore these relationships further in the following section.

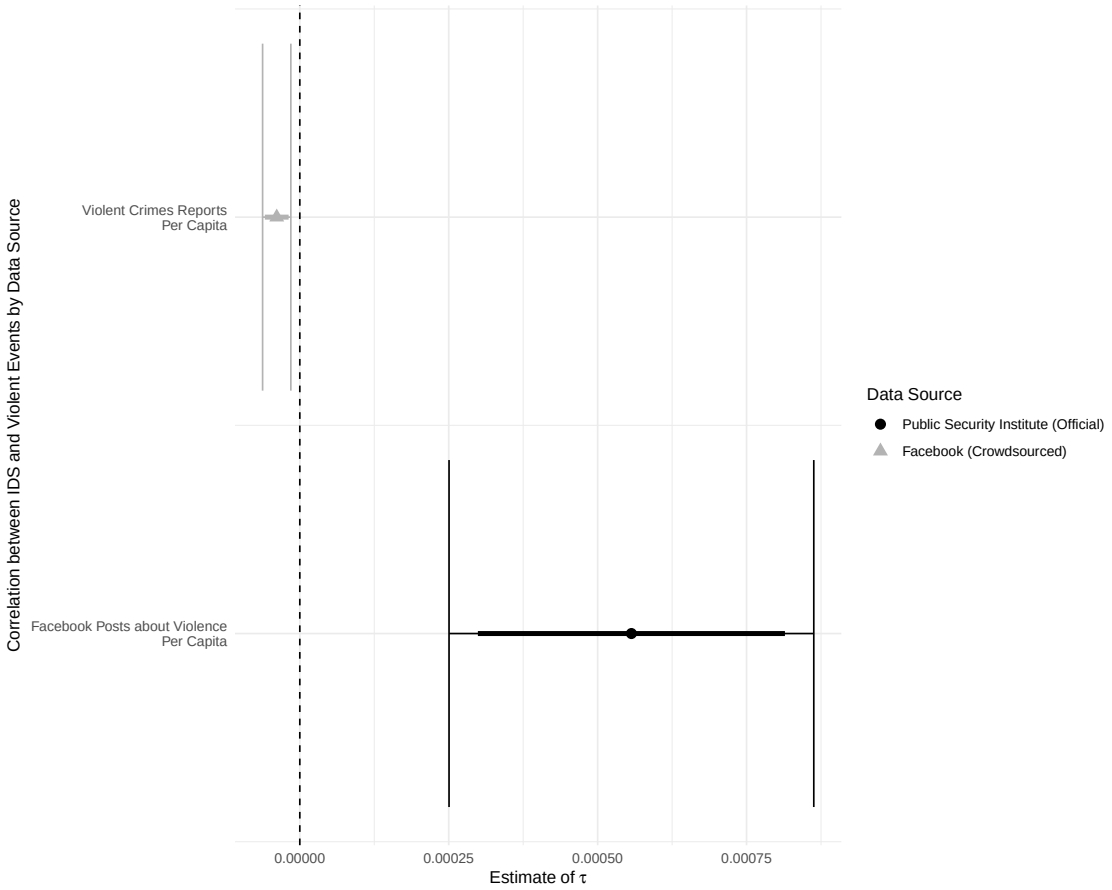

Figure 5.5: Relationship between Social Development, Facebook Posts about Violence and Violent Crime - The figure shows the relationship between violent crime reports per capita and social development (top) and Facebook posts about violence per capita and social development. All models use month-year fixed effects and control for average age.

5.5 Relationship between Reactions to Posts about Violence and IDS

Although official data likely suffers from reporting biases, violent crime in Rio is significantly more common in poorer communities; however, on Facebook pages, I find that an increase in social development also leads to a corresponding increase in posts about violence. This suggests an over-saturation of information about violence in neighborhoods with less actual exposure to violent events. In this section, I examine how neighborhood features such as social development shape content about and reactions to posts about violence on Facebook.

In line with **H3**, I find that social development plays a significant role in shaping citizens' online behavior in community groups. However, it is not citizens in wealthier communities (where violence is much rarer) that are more likely to post about violence as a share of their total posts, but citizens who reside in neighborhoods in the middle development level, who have some exposure to violence, but far less than poorer communities. The proportion of posts about violence was 4.6% higher in community pages focused on neighborhoods in the middle third for development. The proportion of posts about violence was 3% lower in community pages focused on neighborhoods in the highest development levels. These findings suggest that it is the citizens who live in the middle levels of development who disproportionately share information about violence in their communities.

These findings are detailed in **Table 5.2**, which shows the relationship between the proportion of posts about violence per neighborhood and development levels. In model (1) development is measured as a continuous variable. In model (2), I create a dummy variable that divides neighborhoods into three equally sized categories by development levels. Both of these models control for monthly violence, age and population by neighborhood. I use month-year fixed effects.

These patterns are also consistent when looking at how citizens respond to posts about violence in their communities (**H6**). In the middle-development communities, citizens are

Table 5.2: Relationship between Neighborhood Development and Posts about Violence in Community Pages

	Dependent variable:	
	Proportion of Posts about Violence	
	(1)	(2)
IDS	0.123*	
	(0.073)	
Middle Third IDS		0.046***
	(0.009)	(0.009)
Upper Third IDS		0.008
		(0.012)
Constant	0.020	0.179
	(0.093)	(0.062)
Observations	1,957	1,957
R^2	0.104	0.115
Adjusted R^2	0.062	0.074

Table 5.2 shows the relationship between the proportion of post about violence per neighborhood and development levels. In model (1) development is measured as a continuous variable. In model (2) it is measured as a categorical variable with a dummy variable that divides neighborhoods into three equally sized categories by socioeconomic development. All models control for monthly violence, age and population by neighborhood. I use month-year fixed effects. Significance: *p<0.1; **p<0.05; ***p<0.01.

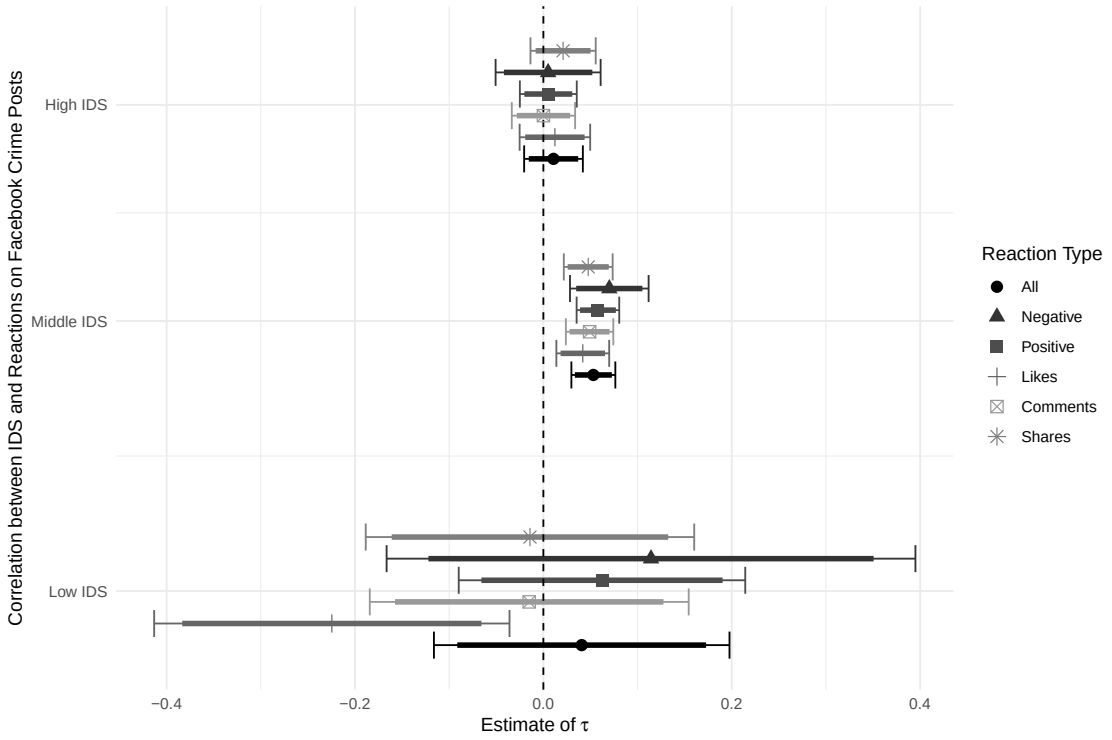

Figure 5.6: Relationship between the IDS and the Proportion of Reactions to Posts about Violence - This figure plots the relationship between social development and reactions to posts about violence as a proportion of total reactions. The figure plots all reactions, as well as divides reactions by type. Negative reactions combine "sad" and "angry." Positive reactions combine "haha" and "love," all of which are ways for users to engage with content on Facebook.

significantly more likely to react to posts about violence across the board. In the least developed and highest developed communities, citizens are equally likely to react to posts about violence, as they are to other types of posts, though there is a wider range among low IDS communities (**Figure 5.6**).

Experiences with Violence and Social Media

What is it about communities with middle levels of development that leads citizens to focus their online attention on violence? As is evident in **Figure 5.7**, In 2019, homicides significantly declined as social development increased. Yet communication about violence online did not accurately reflect this trend, particularly in community groups that are concentrated in the middle with respect to social development. The graph also show much more variation among neighborhoods within the lower and middle thirds than within the top one, suggesting that at least in mid-development level communities there is at least some evidence of violence.

Unlike less developed communities, neighborhoods in the mid-development levels may not try to reclaim the narrative of their environment by sharing stories that document positive growth or change. Furthermore, past research focused on the threat of refugee resettlement has found that negative attitudes are concentrated among those who are less exposed to the perceived threat (Ferwerda, Flynn and Horiuchi, 2017). In the context of Rio de Janeiro, violence is a known quantity in some communities, making each incident less salient or seemingly "newsworthy." This could be another reason why residents in less developed communities, where violence is higher, are significantly less likely to engage with or post about violence in their neighborhoods.

However, unlike more highly developed neighborhoods, residents in mid-development communities are still exposed to at least some violence, which is fairly uncommon in wealthier communities. This is evident in both the official government homicide data (**Figure 5.7**) and through self-reported experiences with violence and victimization, as reported in my survey.

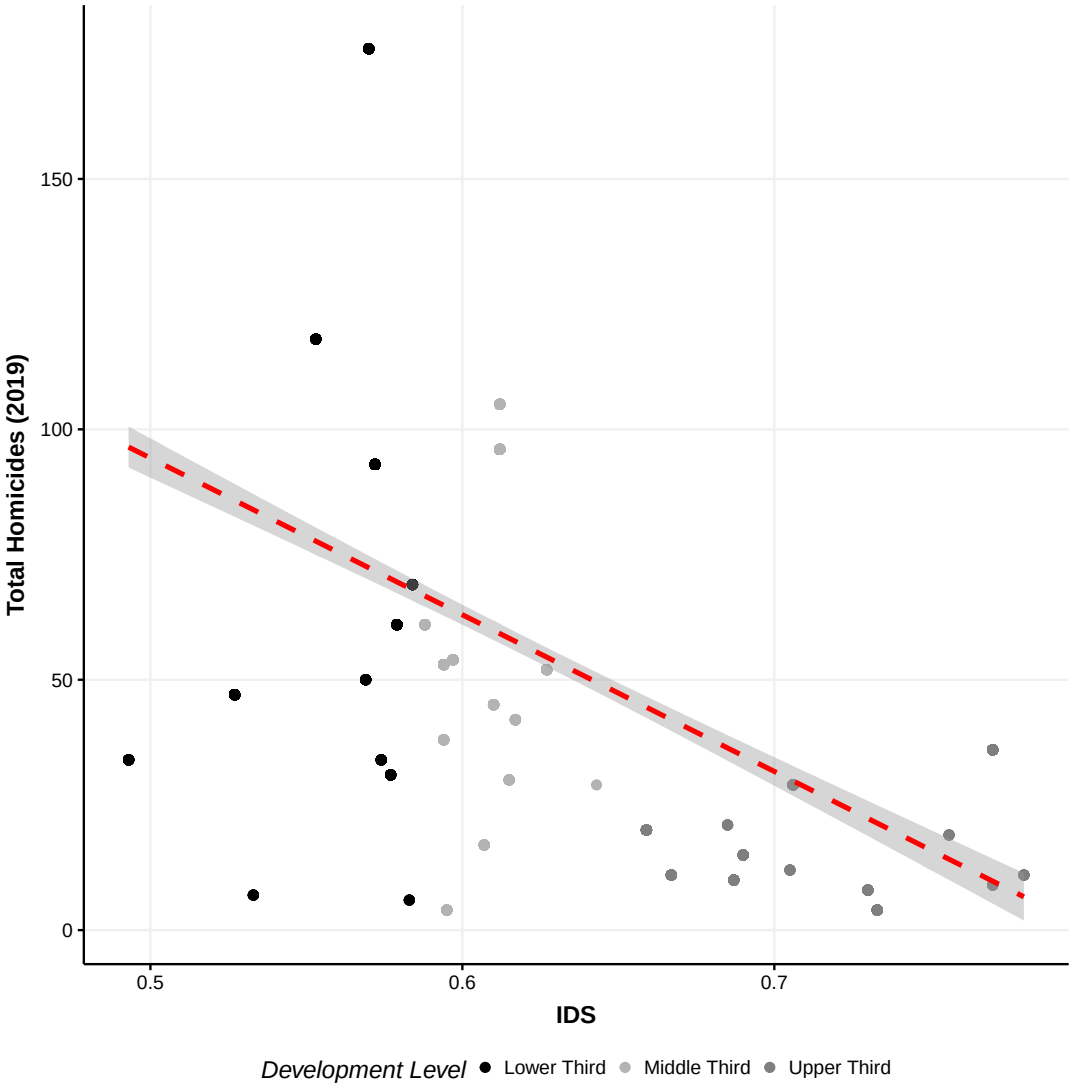

Figure 5.7: Total Homicides (2019), by Development Levels. - Figure shows the relationship between the number of homicides for neighborhoods in Rio de Janeiro in 2019 and development. The color of the points represents the "development level" of each neighborhood as categorized throughout this chapter. Data on homicides come from the Public Security Institute.

Drawing on data from a city-wide, representative survey of Rio de Janeiro conducted in November-December 2020, I find that citizens in mid-development communities were less likely than citizens in less developed communities but more likely than wealthier citizens to report being victimized by narcotraffickers or the police (**Figure 5.8**).

Figure 5.8: Victimization by Development Levels

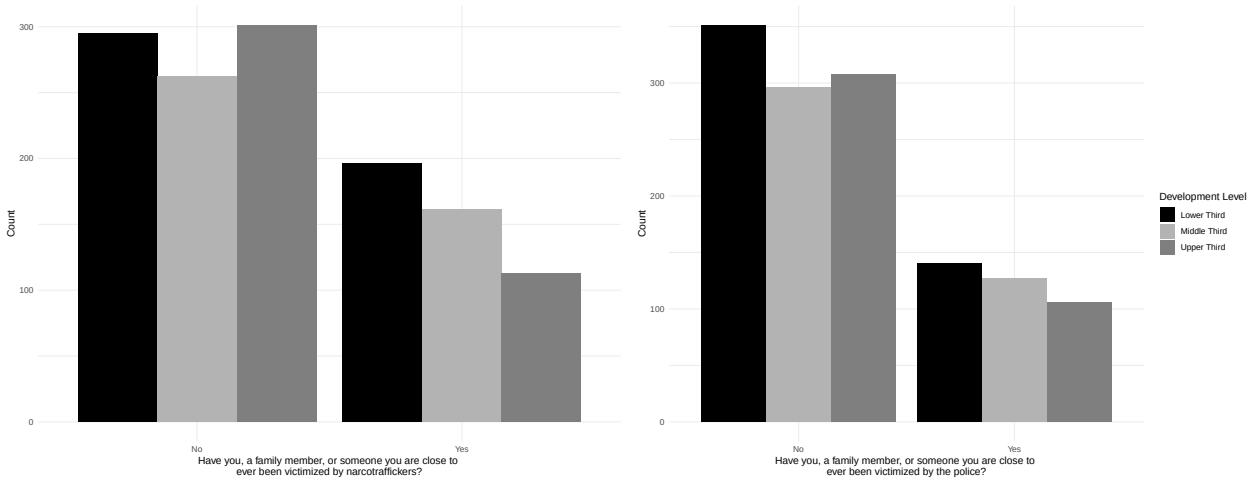

(a) By Narcotraffickers **(b)** By the Police

Figures show survey responses to questions about victimization by (a) narcotraffickers and (b) the police, by development levels.

Although the threat of violence in their community is certainly prevalent, it is less than what citizens in less developed communities face. This makes the proportion of and engagement with posts about violence on mid-development community pages all the more puzzling. If this were a simple story of "shock value" drawing attention, we would expect to see the same rates of sharing and engagement in wealthier communities. What seems more likely is that this phenomenon only works if the perceived threat or fear-evoking incident can *plausibly* be seen as real. A large body of neuroscience and psychology literature finds that the brain – referred to as an "anticipation machine" – utilizes past experiences and information about current conditions to anticipate future challenges (Grupe and Nitschke, 2013; McLaughlin and Hatzenbuehler, 2009).

Citizens in middle income neighborhoods have at least some past experience, which may

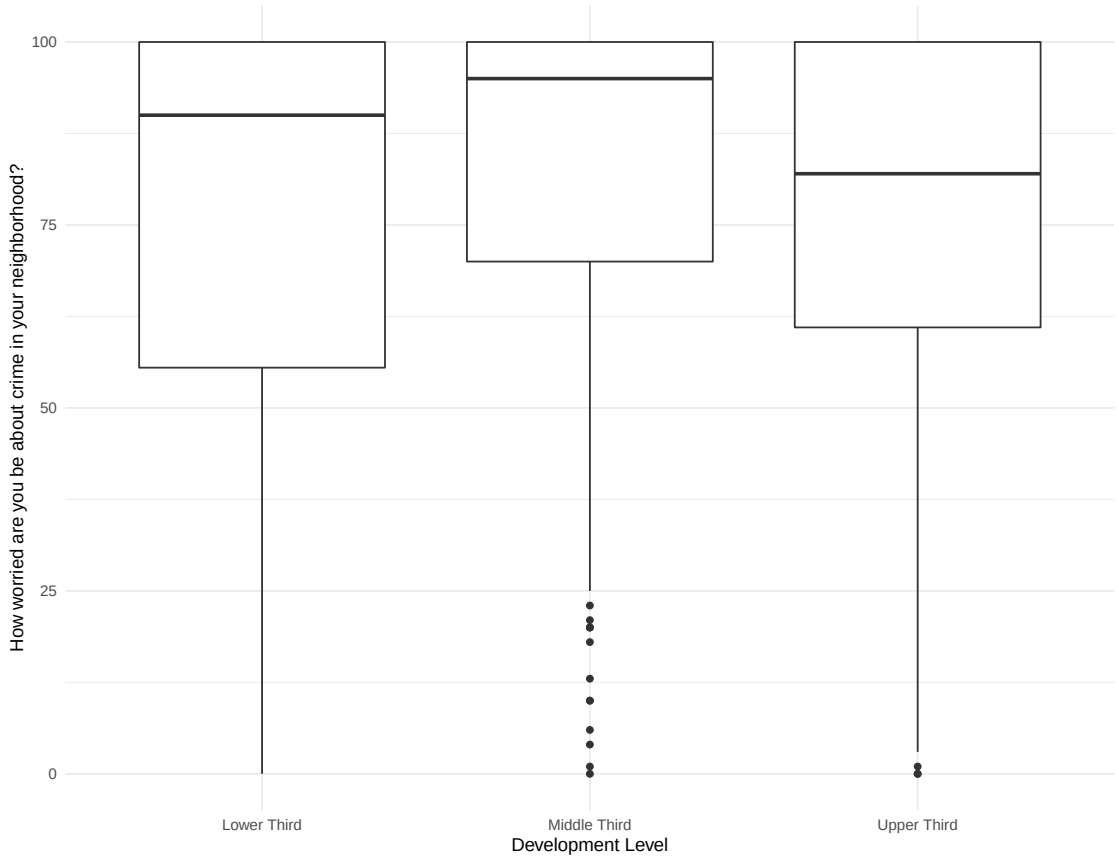

Figure 5.9: Fears about Violence, by Development Levels. - Figure show the interquartile range for two questions on my survey about (a) support for increased policing and (b) concerns about violence, by development levels

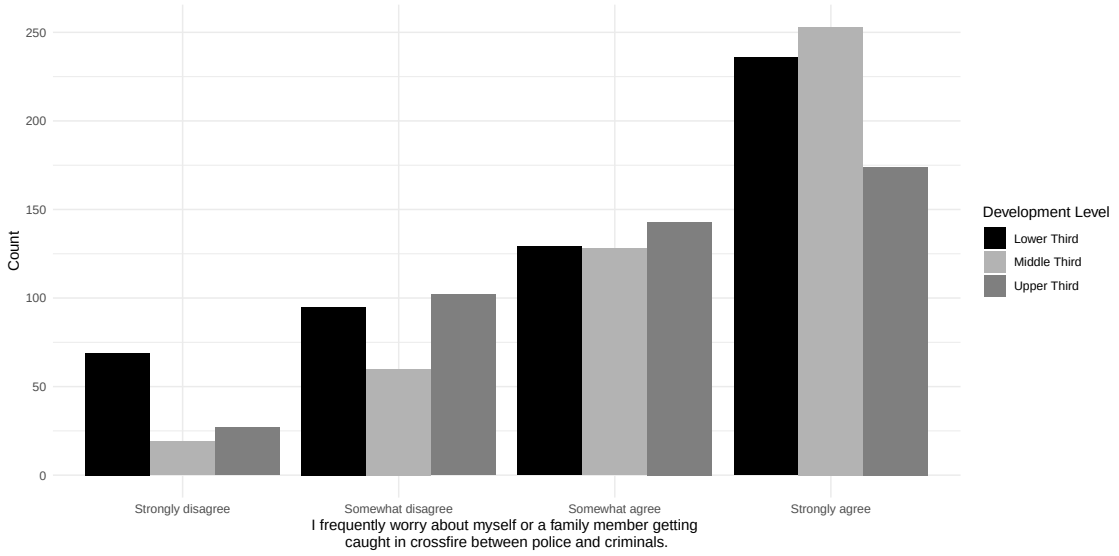

Figure 5.10: Fears about Being Caught in Crossfire, by Development Levels. - Figure show the distribution of responses to a question in my survey about fears of being caught in crossfire, by development levels.

help to shape their future concerns. Drawing on data from my survey, I find that respondents in middle income neighborhoods are significantly more likely than respondents from wealthier neighborhoods to worry about crime in their communities (**Figure 5.9**) and to worry about being caught in the crossfire between the police and criminals (**Figure 5.10**). It is important to note that these patterns are based on preliminary evidence and "snap-shot" observational data; future research should seek to evaluate these dynamics more systematically. However, my findings do suggest that anxiety about security issues is more common among respondents in mid-development neighborhoods than it is for respondents in wealthier neighborhoods, which motivate social media engagement. This engagement stems, at least in part, from past experiences, which – while fewer than in poorer neighborhoods, is not entirely uncommon – motivate concerns about future violence. This pattern reflects research in sociology that describes how despite falling rates of violence in Brazil, feelings about insecurity and fears about violence have increased over time (Borges, 2011). Meanwhile, although actual exposure to violence is higher in less developed neighborhoods, anxiety about violence among respondents in poorer neighborhoods is lower than it is among respondents in middle-income neighborhoods, who may be less fearful of or desensitized to each individual violent event (Funk et al., 2004).

5.6 Conclusion

Drawing on nearly two million Facebook posts from community pages and major media outlets, I evaluate the relationship between the source of content and crime-related communications and community engagement. I find that Facebook users are significantly more likely to react to posts about violence in community pages, and that these pages are significantly more likely to share content about violence as a proportion of their total content. Although this finding challenges perceptions of the media as promoting negative news designed to insight fear – while keeping the audience's attention and driving business – this

is not particularly surprising given the nature of the content. Community pages are particularly well-suited to sharing information about local violence, given their hyper-localized nature.

I further explore these relationships in order to evaluate whether or not online conversations on community pages accurately reflect on-the-ground incidents of violence. I find that while an increase in social development significantly decreases the number of violent incidents, it also increases the number of community posts about violence. This pattern suggests a complicated relationship between the spread of information about violence and actual incidents of it. In evaluating where these patterns diverge, I find that it is communities in the mid-development levels that are significantly more likely to share and react to posts about violence online. I argue that this divergence occurs for two reasons: (1) less developed communities may be somewhat desensitized to individual violent events and less interested in sharing content that paints their communities in a negative light; and (2) as compared to wealthier communities, middle income communities have at least some experience with violence, which colors their concerns about future violent events. I provide evidence for these conclusions using data from an original survey of 1,245 respondents.

These findings are important for a number of reasons. First, I systematically document a clear way that social media content differs from the mainstream media and note that it may be even more prone to sensationalism than more traditional sources of information. I also demonstrate that social media is prone to biases that may not accurately represent the on-the-ground reality. This is particularly important given that the user serves as both the content creator and curator on a number of social media platforms, thereby introducing their own biases into the information that they generate and share. Because citizens are more likely to trust content if it is created by people "like them," they will also be more exposed to information that is biased in some way. With respect to violence, these biases may include heightened fear about crime, which then colors the content that is created – thereby raising overall perceptions of violence in a way that may not accurately reflect reality. In these

contexts, such perceptions may be enough to significantly impact policy decisions and shape support for positions driven by heightened anxieties about violence.

6

Conclusion

In high-income countries the adoption of social media has largely plateaued; however, in many low- and middle-income countries, engagement on these platforms has continued to skyrocket, emerging as an increasingly important way for citizens to acquire information about the world (Poushter, Bishop and Chwe, 2018). Despite this trend, there remains a limited body of research devoted to these contexts, that examines the ways in which information obtained through social media may differ from other sources and how these differences may interact with the political environment to shape public opinion and subsequent policy preferences. In particular, the low barrier to entry and the reliance on user-generated content distinguish social media platforms from other information mediums. In my research, I find that these factors have implications for both information processing and attitude formation, as well as the accuracy of the information environment.

The research presented herein fills a theoretical gap in both the literature on public opinion formation and the literature on the political implications of social media usage. It does so by exploring the effect of information provided through social media – particularly about salient topics such as violence – on political attitudes and perceptions of insecurity among voters in Brazil. Specifically, I ask: (1) does the citizen-driven nature of information from social media alter voters' political beliefs more than information provided by more

traditional gatekeepers?; and (2) does the effect of this information differ based on the uncertainty of prior beliefs among those who receive it? In this chapter, I summarize what this research finds about the political effects of social media usage in Brazil, in particular. I also highlight its potential applications in low-and middle-income countries around the world more broadly.

6.1 The Political Effects of Crowdsourcing Information about Violence

Uniting insights from public opinion and political communications research, I develop a theory of information processing in the social media age that incorporates both (a) the uncertainty of individual attitudes and (b) the user-generated nature of that information, which relies on "the crowd" as content creators and curators to produce a seemingly unlimited amount of information. Building on extant literature (e.g., Arias et al., 2018), I argue that, in general, higher uncertainty about prior beliefs leads to greater information processing; however, social media – due to both the role of everyday citizens in creating content and curating exposure – is distinct from other information sources (e.g., Farrell, 2012). These distinctions may be particularly appealing in contexts of low government trust or low media capacity – contexts where social media usage is rapidly growing. Citizens in these places may view information provided by "people like them" as more reliable and thus be more willing to incorporate it into their attitudes or beliefs. I expect this to be the case particularly among individuals who are uncertain about their beliefs, due to the fact that, in a Bayesian sense, the range of values across which the truth can fall is more uniform and uncertain individuals may be less likely to dismiss evidence that contradicts their strongly-held, preexisting beliefs.

Just because people may be more willing to update their beliefs based on new information from social media does not mean that their overall perceptions become more informed or more closely resemble the truth. Social media by definition relies on everyday users to create

content, thereby eliminating the need for editors or media executives (Shirky, 2008). Due to the biases of individual users, who may be more likely to share shocking or salient, if uncommon, incidents, the content created through these channels may exaggerate certain types of events (e.g., a rare violent incident), while ignoring incidents that are far more common.

I explore the relationship between social media information and public opinion using a mixed methods approach that combines experimental evidence and observational data. I evaluate the effect of citizen-to-citizen communication via social media and examine if and how it alters political attitudes and perceptions of insecurity. To measure uncertainty, I develop an index ranging from 0 to 1 that is based on respondents' self-reported uncertainty with respect to five questions about violence. I then randomized participants to receive one of two information prompts about neighborhood violence – one included standard information from the government (control) and the other included crowd-sourced information from a social media platform (treatment).

I find that when compared to information from government sources, information from social media platforms significantly alters concerns about crime (Outcome 1) and support for the police (Outcome 2). However, I find that the directions of these effects are conditional on levels of uncertainty. For both outcomes, low uncertainty individuals who are exposed to social media information are more likely to harden their existing beliefs, while high uncertainty individuals are more likely to reverse course. Regardless of uncertainty levels, outcomes do not significantly change among individuals who receive information about violence from the government.

I also explore the mechanisms underlying these findings. In particular, I focus on the role of trust in shaping changes in attitudes. Using a conjoint experiment, I ask respondents to choose which of two randomly generated information profiles they would trust more. I vary the following attributes: (1) source, or specifically who is providing this information; (2) engagement, or how many users have boosted this content by engaging with it; (3) timing,

or when was the information shared; and (4) associated evidence, or does it have additional photos or videos. I also vary whether the information they receive is about violence or government services. I find that regardless of the focus of the information, citizens are more likely to trust information if it comes from sources similar to them (e.g., not from the government), has higher levels of user engagement, is posted in real time, and has associated evidence. These attributes are all features of social media information, suggesting the vital role of trust in shaping the willingness to believe content.

In the final empirical chapter of my dissertation, I ask: if citizens are more likely to trust content they see on social media, does that content reflect the same reality presented through more official channels? In theory, more information should lead to a more informed electorate by boosting the accountability required for a representative government. This theoretical expectation may not play out as anticipated, however, if the information is flawed or biased in important ways. Drawing on hundreds of thousands of social media posts from Facebook, I examine the patterns of information dissemination across Rio de Janeiro and evaluate biases in crime-related communications relative to baseline crime rates and other neighborhood characteristics. I find a clear disconnect between mainstream media and government data, on the one hand, and social media content about violence, on the other. Specifically, social media pages are significantly more likely to share content about violence when compared to mainstream media content. This is not particularly surprising given that community pages on Facebook may be better suited to sharing local level information about violence than mainstream media outlets, which due to capacity constraints, cannot treat every incident of violence as newsworthy.

To further explore how the spread of information on social media may differ from other information purveyors, I evaluate the relationship between government data about violence and the information shared on community pages. I find that community pages that serve less developed communities are far less likely to share content about violence on social media, but far more likely to experience incidents of violences, as reported in the neighborhood

government data. Looking closer at the content on community pages, I find that content about violence is specifically overrepresented on community pages that serve neighborhoods in the mid-development range. A post hoc explanation is that wealthier communities may have very limited amounts of violence, and thus the content does not register on community pages. In mid-development communities, there is some violence to be discussed, and each violent event may generate more conversation. In low-development communities, where violence is more common, community pages may serve a different purpose. Rather than share information about violence, they seek to challenge narratives about these communities and specifically highlight positive aspects of community life. These findings could also be due to systematic biases in government data, which underreport incidents in more developed communities. To fully understand the differences between crowdsourced information and government information about violence will require further investigation.

6.2 Further Applications

Much of what we know about how social media shapes the information environment and subsequently alters political outcomes stems from research conducted in the United States. Although this research provides critical insights into the effect of social media on public opinion, the United States is unique for a variety of reasons, including the nature of electoral politics, the expansive media market, and the foundational precedent of electoral democracy. As a result, general theories of change established in this context may not be exportable outside of the United States or to countries that may be emerging from conflict or dictatorial regimes, for example. Importantly, in contexts where the use of social media is continuing to grow, understanding the political impact of these platforms is all the more valuable due to the theoretical link between an informed public and the establishment of democratic norms.

In this dissertation, I test the theory detailed in **Chapter 3** in a Brazilian context. This is because from a development, press freedom, and democratic governance perspective,

Brazil is noticeably average when compared across countries. Not long ago, the Brazilian government was also run as a military dictatorship, making the transition to democracy a relatively recent phenomenon.[1] Furthermore, social media is widely used across Brazil, and some have already criticized it for its role in shaping electoral outcomes.

Due to the size and geographic importance of Brazil, understanding these dynamics in this context is valuable in its own right. However, I would also expect this theory, and the empirical implications that follow, to be applicable in a variety of contexts. Specifically, I would expect aspects of the theory to apply to young democracies, emerging from conflict settings or authoritarianism. These countries often share a general skepticism of the government and have less-developed media markets. At the same time, they do have elections, which may make the accumulation of information for public opinion more valuable. Finally, although I would expect similar dynamics to play out in non-democratic countries, or countries with ongoing conflicts, additional variables – for example, Internet censorship – may complicate these expectations and are beyond the scope of this project. Future research, however, should explore both of these conditions in greater detail.

6.3 Contribution and Implications

This dissertation represents an early effort to evaluate the ways in which public opinion formation is shifting as a result of social media information outside of a U.S. context. From a theoretical perspective, I unite and build on related literature in disparate fields to develop a novel theory of public opinion formation in the social media age, building on insights from political science, psychology, communications, and economics. I test this theory using an original survey, which incorporates several survey experiments and details a novel way to measure uncertainty. I also utilize novel Natural Language Processing tools to evaluate a large number of text-based observations.

[1]The dictatorship in Brazil lasted 21 years and formally ended in 1985. The first direct presidential election was held in 1989.

The focus here on crime and violence is also timely, given that it is one of the more common types of information shared by citizens on social media platforms. As social media continues to play a vital role in information dissemination, this type of content can both shed light on previously neglected violent events and be used for other more nefarious objectives, such as racial profiling, disseminating disinformation, or overdramatizing violence. For this reason, it is important to understand how individuals process and react to information from social media platforms. This is particularly true in contexts such as Rio de Janeiro, where social media is deeply entrenched in daily life due low state capacity, but also critical in places such as the United States where online platforms have played an important role in exposing police brutality yet may also be weaponized for other ends. Disentangling these processes – and understanding how they shape voter behavior and perceptions of insecurity – is particularly important as we continue to grapple with the influential, information disseminating power of social media platforms.

7

Supplementary Materials

7.1 Chapter 4

Survey Text for Survey Experiments

Demographics Questions

1) In which community do you live (or live closest to)? Please select the best available option.

- Bangu
- Barra da Tijuca
- Bonsucesso
- Botafogo
- Brás de Pina
- Campo Grande
- Catete
- Cidade Nova
- Copacabana
- Engenho Novo

- Gávea
- Guaratiba
- Honório Gurgel
- Ilha do Governador
- Inhaúma
- Ipanema
- Leblon
- Madureira
- Marechal Hermes
- Méier
- Mem de Sá
- Pavuna
- Penha
- Piedade
- Praça da Bandeira
- Praça da República
- Praça Mauá
- Praça Seca
- Realango
- Recreio dos Bandeirantes
- Ricardo de Albuquerque
- Rocinha
- Santa Cruz
- Santa Teresa
- São Cristóvão
- Tanque
- Taquara

- Tijuca
- Todos os Santos
- Vicente de Carvalho
- Vila Isabel

2) What is your age?

- 18-25
- 26-40
- 41-60
- 61+

3) What is the highest level of education you have received?

- Less than Secondary School
- Secondary School Degree
- Some University but no degree
- Bachelor's Degree in University (4-year)
- Graduate degree (MA, Ph.D, JD, MD, MBA)

4) What is your sex?

- Male
- Female

5) What is the total monthly income of all people living in your household?

- Up to R$1.045,00 (up to 1 minimum wage)
- From R$1.046,00 to R$2.091,00 (more than 1 and up to 2 minimum wages)
- From R$2.092,00 to R$5.226,00 (more than 2 and up to 5 minimum wages)
- From R$5.227,00 to R$10.451,00 (more than 5 and up to 10 minimum wages)
- From R$10.452,00 to R$15.675,00 (more than 10 and up to 15 minimum wages)

- From R$15.676,00 to R$20.900,00 (more than 15 and up to 20 minimum wages)
- More than R$20.901,00 (more than 20 minimum wages)

6) Which political party do you identify with most?

- PSL
- PT
- PSDB
- PSD
- PP
- MDB
- PSB
- PL
- PRB
- DEM
- PDT
- PSOL
- NOVO
- PODE
- PROS
- PTB
- PRTB
- PCdoB

7) With respect to race, which category do you identify with?

- White
- Pardo
- Black
- Indigenous

- Other

- I don't know/I prefer not to say

Conjoint Experiment

Four times each, randomly ordered mix of development project and local violence prompt.

(8)-(11) You see a post on Facebook about an announcement related to government service provision in your neighborhood. It had the following features: (features). Which of the options presented would you be more likely to believe?

(12) - (15) You see a post on Facebook about an incident related to violence in your neighborhood. It had the following features: (features). Which of the options presented would you be more likely to believe?

Randomize which levels from each factor the respondent will see in two side-by-side pairings. The respondent will then be asked to choose which profile they would find more credible.

Factor	Levels
Timing	(1) *The information was shared at the time of the incident*; (2) the information was shared a day later
Source	The person who posted the story: (1) *was a member of your neighborhood*; (2) was a city councilor; (3) was a member of the media; (4) *was someone you know personally*
Engagement	(1) *The post had more than 500 likes, shares, and comments*; (2) *The post had around 100 likes, shares, and comments*; (3) The post had less than 20 likes, shares, and comments
Media	(1) The story included photos of the incident; (2) The story included video of the incident; (3) the story included neither photo nor video

General Questions

16) Please rank the following issues based on your priorities when voting.

- Crime prevention
- Local development
- Corruption prevention
- Healthcare
- Education
- Unemployment

17) Are you part of a Facebook group or page focused on your neighborhood?

- Yes
- No

18) Please rank the following source from most used to least used for information about your neighborhood.

- Television
- Newspapers
- Facebook
- Radio
- WhatsApp
- Twitter
- Your neighbors, friends or relatives

19) How often do you get information about your neighborhood from Facebook?

- Often
- Sometimes
- Rarely

- Never

20) How often do you get information about your neighborhood from WhatsApp?

- Often
- Sometimes
- Rarely
- Never

21) How often do you get information about your neighborhood from Twitter?

- Often
- Sometimes
- Rarely
- Never

Please read the following sentences and indicate your agreement or disagreement.

22) I am more concerned about the future than I was a year ago.

- Strongly disagree
- Somewhat disagree
- Somewhat agree
- Strongly agree

23) My neighborhood is worse off economically than it was a year ago.

- Strongly disagree
- Somewhat disagree
- Somewhat agree
- Strongly agree

24) The police are a positive force in my neighborhood.

- Strongly disagree

- Somewhat disagree
- Somewhat agree
- Strongly agree

25) I frequently worry about myself or a family member getting caught in cross-fire between police and criminals.

- Strongly disagree
- Somewhat disagree
- Somewhat agree
- Strongly agree

40) How worried are you about urban crime in your neighborhood, ranging from 0-100, with 0 being "not at all worried" and 100 being "extremely worried."

[number from 0-100]

41) How likely are you to support increased police presence in your neighborhood, ranging from 0-100, with 0 being "extremely unsupportive" and 100 being "extremely supportive."

[number from 0-100]

Knowledge and Experiences Questions

Please answer the following questions to the best of your ability.

27) Approximately how many homicides were there in your neighborhood last year?

(input number)

28) How certain are you about your response to the previous question?

- Very certain
- Somewhat certain
- Somewhat uncertain

- Very uncertain

29) Has violence gone down in your neighborhood since 2019?

- Yes
- No

30) How certain are you about your response to the previous question?

- Very certain
- Somewhat certain
- Somewhat uncertain
- Very uncertain

31) Which police force is responsible for investigating crime?

- Military Police
- Civil Police
- Pacifying Police Force
- Special Police Operations Battalion

32) How certain are you about your response to the previous question?

- Very certain
- Somewhat certain
- Somewhat uncertain
- Very uncertain

33) On which street is the military police battalion in your neighborhood located?

(input street)

34) How certain are you about your response to the previous question?

- Very certain
- Somewhat certain

- Somewhat uncertain
- Very uncertain

35) What is the general telephone number of the Military Police?

(input number)

36) How certain are you about your response to the previous question?

- Very certain
- Somewhat certain
- Somewhat uncertain
- Very uncertain

37) Have you, a family member, or someone you are close to ever been victimized by the police?

- Yes
- No

38) Have you, a family member, or someone you are close to ever been victimized by narcotraffickers?

- Yes
- No

Information Experiment

Participants will receive one of the two pieces of information, and then respond to the same questions.

According to the Institute for Public Security, which provides government data about violence in Rio de Janeiro, the civil police reported [NUMBER] homicides in your neighborhood of [NEIGHBORHOOD] between January and August of this year. These incidents were reported online a month later.

OR

According to a local website which provides citizen-to-citizen information on neighborhood gunshots, residents reported [NUMBER] violent incidents in your neighborhood of [NEIGHBORHOOD] between January and August of this year. These incidents were reported in real-time, while the violence was occurring.

39) Is this lower or higher than you expected?

- Much lower
- Somewhat lower
- About the same
- Somewhat higher
- Much higher

39) How likely are you to believe this information, ranging from 0-100, with 0 being "extremely unlikely" and 100 being "extremely likely."

[number from 0-100]

40) How worried are you about urban crime in your neighborhood, ranging from 0-100, with 0 being "not at all worried" and 100 being "extremely worried."

[number from 0-100]

41) How likely would you be to support increased police presence in your neighborhood, ranging from 0-100, with 0 being "extremely unsupportive" and 100 being "extremely supportive."

[number from 0-100]

De-Brief: *Thank you for your participation. Please note that this research is not affiliated with any of the institutes or companies referenced and all scenarios described in this survey are fictitious and do not reflect actual events. If you have any questions or comments please feel free to send them via email to vwirtschafter@ucla.edu.*

7.2 Chapter 5

Table 7.1: Rio de Janeiro Community Pages

No.	Group Name	Data Created	Followers (2020)	# Posts
1	A Voz de Rio das Pedras	5/10/13	14,435	1,180
2	A Voz do Lins de Vasconcelos	10/21/13	9427	1,411
3	Abolição Piedade DEPRESSÃO	3/5/13	33,570	995
4	Alerta Abolição Piedade e Pilares	4/3/17	8,064	475
5	Alerta Noticias Da Baixada	7/17/18	17,254	1
6	Alerta Barra	3/14/17	40568	719
7	Alerta Barra da Tijuca	10/22/12	14,178	267
8	Alerta Flamengo	9/17/16	2,769	575
9	Alerta Geral Tijucano	7/2/15	25,669	495
10	Alerta Glria e Catete	3/27/19	2,152	140
11	Alerta Guadalupe NL	9/25/18	152,133	1248
12	Alerta ILHA 24h	9/17/17	44,054	3958
13	Alerta Informes RIO	9/23/17	25,170	215
14	Alerta Jacarepagua	2/3/18	2,146	39
15	Alerta Jardim Oceânico	11/17/16	18,889	893
16	Alerta Largo do Bicão Notícias	1/23/15	12,682	1,583
17	Alerta Recreio	12/6/16	14830	245
18	Alerta Vila Isabel Original	5/15/15	65,944	2,917
19	Alertando Geral	1/29/13	23,836	3,530
20	Alto da Boa Vista-RJ	9/25/10	4,392	384
21	Amigos de Anchieta	7/26/14	8,097	3,024
22	Amigos Do Bairro De Quintino Bocaiva	8/20/15	2,285	810
23	Anchieta Agora RJ	2/19/17	23,994	5,879
24	Anchieta Em Alerta	10/20/17	26,295	4
25	Anchieta News	3/28/14	9,904	66
26	Aqui em Curicica	3/11/16	17,345	1,680
27	Assaltos em VAZ LOBO	12/8/14	13,502	444
28	Bairro de cascadura	8/13/14	7,490	1,426
29	Bairro de Vista Alegre - Rio de Janeiro/RJ - Brasil	2/18/12	13,424	7,352
30	Bairro Maria da Graßa	4/19/16	5,552	1,040
31	Bangu News	1/15/13	18,693	859
32	Bangu News Ao Vivo	12/8/15	8,878	20
33	Bangu News Crimes	3/24/15	72,969	985
34	BANGU Urgente	5/16/16	11,684	4,156
35	Barra da Tijuca da Depressão	10/5/12	9,354	16
36	Barra da Tijuca News	9/13/16	18,011	404

No.	Group Name	Data Created	Followers (2020)	# Posts
38	Barra de Guaratiba.rj	7/14/16	14,221	166
39	BENTO Ribeiro News	9/17/15	13,492	566
40	Bento Ribeiro News	7/21/15	11,725	933
41	Blog Crimes News	9/24/15	195,075	1,409
42	Blog do Jacaré - Notícias de Jacarepaguá	3/22/12	28,041	1,612
43	Bonsucesso da Depressão	11/1/12	51,051	1,867
44	Bota Cara CDD	2/1/16	13,450	303
45	Brás de Pina News	9/20/17	7,430	323
46	Campo Grande Notícias	1/23/18	57,576	307
47	CANAL DO ANIL	11/4/12	51,620	5,330
48	Caos No Rj	9/13/17	24,295	1,363
49	Cascadura - Caminhos do Subúrbio	5/23/12	26,843	11,328
50	Cascadura News	5/10/16	30,041	4,776
51	Catumbi da Depressão	1/16/13	9,242	693
52	Catumbi Urgente	8/29/16	3,189	51
53	Cdd Acontece	8/2/11	122,510	6,010
54	Central Rio Das Pedras	2/3/14	35,048	1,404
55	Centro/lapa Residencial	3/2/17	6,752	2,140
56	Chatuba News	5/21/15	19,793	514
57	Coletivo Fala Akari	11/15/15	18,521	966
58	Coletivo Papo Reto	4/23/14	54,947	3,871
59	Comunidade do Jacarezinho	10/31/13	5,931	48
60	Copacabana News	2/20/14	2,533	159
61	Cordovil News	11/13/11	8,711	109
62	Costa Barros - RJ	10/1/18	38,880	729
63	Curicica	9/21/11	10,810	1,770
64	De Olho na Notícia	11/6/18	15,191	611
65	Deodoro da Depressão	4/11/13	3,220	244
66	Deodoro Informa	2/19/16	2,178	316
67	Eng.Novo Lins, Méier Informes	6/6/15	20,906	2,167
68	Engenheiro Leal News	7/10/17	3,195	614
69	EU MORO EM BANGU	9/21/12	80,094	36,297
70	Fala Manguinhos	5/22/13	24,868	3,237
71	Fala, Abolição	8/22/16	3,758	343
72	Favela Fiscal	3/27/13	4,360	1,792
73	Favela Não Se Cala	11/25/14	36,992	1,328
74	Favela Santa Marta	9/21/10	6,283	586
75	Fazenda Botafogo - Bairro	12/17/13	10,021	455
76	Fazenda Botafogo - RJ	9/25/11	4,060	148
77	Fazenda Botafogo Depressão	9/4/12	3,688	37
78	Fogo Cruzado RJ	7/4/16	74,825	18,366
79	Freguesia da Depressão	1/16/13	47,779	2,333
80	Gardénia Azul Notícias	11/18/16	32,608	2,159
81	Graja RJ	10/3/11	4,329	232
82	Grande Méier da Depressão	9/15/12	101,954	3,199
83	Guadalupe da Depressão	11/11/12	7,222	204
84	Guadalupe News	12/25/13	248,942	7
85	Guanabara Alerta	7/13/14	8,802	201
86	Guerra Urbana RJ	1/23/13	53,937	1,245
87	Honório Gurgel	7/8/13	3,834	237
88	Honório Gurgel News	1/21/17	8,289	508
89	Ilha Notícias	2/5/12	150,737	9,641

No.	Group Name	Data Created	Followers (2020)	# Posts
90	Informe Bangu	1/24/18	28,305	1
91	Informe Coelho Neto	2/26/16	33,148	2,698
92	Informe Irajá	5/31/16	11,277	1,820
93	Informe Madureira News	9/27/16	3,698	1
94	Informe Marechal Hermes RJ	4/27/16	38,868	1,550
95	Informe Rocha Miranda	6/17/16	73,474	5,804
96	Inhaúma da depressão	9/17/12	5,843	269
97	Inhaúma Notícias	9/27/17	19,613	513
98	Inhoaóba News	1/31/17	5,227	81
99	Ipiranga News	6/16/16	44,617	2,798
100	Irajá News RJ	7/14/16	131,923	1,110
101	Irajá Notícias	10/21/18	3,883	12,789
102	Itanhangá Noticias	10/6/19	7,691	231
103	Jacarepaguá alerta Tiros - RJ	2/7/14	416,861	593
104	Jacarepaguá Notícias RJ	6/3/15	302,098	41,062
105	Jacarepaguá Online	2/27/13	28,544	19,833
106	Jornal Alemão Notícias	4/19/15	14,265	2,487
107	Jornal Complexo da Penha Notícias OnLine	5/16/14	3,818	893
108	Jornal da Barra	8/7/14	2,934	1,282
109	Jornal da Rocinha	4/14/16	8,795	69
110	Jornal do Bairro	3/16/13	14,994	1,793
111	Jornal Fala Roça	2/1/11	15,269	971
112	Jornal Nosso Bairro Jacarepaguá	1/15/13	22,166	1,610
113	Jornal O Debate 24 Horas	1/21/16	263,915	4,030
114	Jornal Paz no Rio News	11/19/15	29,079	19,536
115	Jornal Praça Seca News	6/14/11	16,517	2,114
116	Jornal Rocinha Notícias	10/21/12	7,238	112
117	Laranjeiras Online	7/6/17	2,597	525
118	Leblon Alerta	6/6/16	43,078	42
119	Leblon News	8/29/16	5,400	1,046
120	Leme Alerta	10/18/15	34,085	508
121	Lins & Engenho Novo / Fatos e Notícias	3/10/13	15,08	3,608
122	Lins da Informação	2/26/17	7,680	2,249
123	Madureira Alerta RJ	2/8/16	217,948	47
124	Madureira News RJ	1/27/15	6,516	11,121
125	Magalhães Bastos	11/1/19	5,124	2,299
126	Magalhães Bastos RJ notícias	3/28/14	187,680	353
127	Maré Vive	8/8/18	6,934	4,652
128	Marechal Hermes Bairro	2/17/19	16,949	444
129	Marechal Hermes News	5/5/17	12,064	205
130	Marechal News	2/7/14	11,847	328
131	Melhora Marechal Hermes	1/5/13	14,671	1,330
132	Meu Leme Rio de Janeiro	1/10/14	22,863	4,265
133	Minha Penha	11/26/12	41,036	2,127
134	Morar em Santa Cruz Da Depressão	8/16/18	7,519	1,049
135	Notícias de Bangu	11/10/16	17,549	9
136	Notícias de Penha/SC	3/18/18	3,127	1,825
137	Notícias Méier	10/28/14	25676	309
138	Noticias de Campo grande rj	3/27/13	18,518	1,438
139	O Morador é o Centro	6/10/15	105,869	3,166
140	O povo de sepetiba	1/1/14	37,061	6,833
141	Olaria da depressão	1/1/16	743,726	2,717
142	Onde Tem Tiroteio-RJ. OTTRJ	8/15/18	2,670	27,141
143	Paciéncia Online 24 Horas - Notícias	5/1/16	14,764	194

No.	Group Name	Data Created	Followers (2020)	# Posts
144	Paciéncia RJ News	4/6/20	9,189	993
145	Parceiros da Rocinha	11/28/16	30,968	1,708
146	Parceiros do RJ	7/11/17	14,405	849
147	Parceiros do Vidiga	12/20/12	83,710	14,439
148	Pavuna da Depressão	12/11/13	80,924	3,754
149	Pavuna Notícias	12/14/16	7,153	331
150	Pavuna pede socorro	12/16/13	3,087	114
151	Pechincha da depressão	3/24/13	11,593	3
152	PEDRA DE GUARATIBA	4/21/11	24,425	2,927
153	Penha - Cordovil - Cidade Alta - (Rj)	7/17/12	2,588	133
154	Penha da depressão	1/3/13	78,025	5,093
155	Penha News	3/12/17	100,928	12,410
156	Penha News - Rio de Janeiro	4/17/12	31,764	2,983
157	Penha Notícias	3/25/14	10,707	882
158	Plantão RJ	7/24/18	2,783	283
159	Portal de Sepetiba	11/28/14	19,842	1,098
160	Portal Ipanema	3/18/15	17,372	2,236
161	Praça Seca	8/10/10	24,164	1,929
162	Praça Seca da Depressão	1/19/13	9,945	151
163	Rádio Rocinha RJ	9/1/11	8,277	161
164	Ramos da depressão	1/13/14	42,524	2,912
165	Realengo Ao Vivo	12/13/16	7,921	338
166	Realengo News	4/3/18	9,330	1,961
167	Realidade Do RJ	10/11/17	120,012	5,687
168	Recreio Noticias	9/14/16	3,110	125
169	Redes da Maré - Somos Todos Maré	8/3/11	56,640	5,473
170	Ricardo De Albuquerque News	6/6/18	10,865	212
171	Rio Comprido Alerta	4/23/15	118,996	4825
172	Rio da Prata News Bangu	8/11/15	16,595	312
173	Rio Das Pedras	5/9/13	12,870	93
174	Rio de Janeiro - Guerra Civil Urbana	3/26/16	96,216	2,748
175	Rio De Janeiro Notícias	1/12/16	295,219	896
176	Rio de Tiros e Assaltos	9/10/16	26,123	195
177	RJ Informe News	9/19/15	111,345	10,324
178	Rj news	8/23/19	32,573	48
179	RJ Notícias	9/29/16	10,712	402
180	Rocha Miranda da Depressão	7/27/13	12,793	51
181	Rocha Miranda news	9/6/15	7,934	327
182	Rocinha Alerta	1/3/17	66,123	5,400
183	Rocinha em Foco	8/30/13	108,556	14,538
184	ROCINHA.ORG	8/6/11	23,712	419
185	S.O.S Brás de Pina	6/10/13	12,372	530
186	São Cristóvão	2/14/12	64,211	5,736
187	Salvemos São Conrado	8/30/12	41,807	4,490
188	Santíssimo News	9/7/16	36,627	3,052
189	Santíssimo Notícias RJ	9/3/13	73,795	8,463
190	Santissimo RJ News	5/29/15	29,033	1,872
191	Senador Camará Online	3/26/16	22,962	600
192	Sepetiba em Alerta	5/23/13	3,154	317
193	Sepetiba NoticiaRj	4/11/18	16,200	1,710
194	SOS Jardim América	1/10/14	2,352	151
195	SOS Valqueire	6/2/16	127,205	6,817

No.	Group Name	Data Created	Followers (2020)	# Posts
196	Sulacap News	5/21/14	41,672	6,738
197	Tó de olho news	8/24/15	137,217	3,624
198	Tanque/Pechincha da depressão	1/12/13	3,420	221
199	Taquara News	1/15/14	18,382	1,027
200	Tijuquinha - Itanhangá	1/19/11	2,271	24
201	Tomás Coelho da Depressão	1/24/13	24,646	1,638
202	Valores da Penha	11/11/16	5,939	1,764
203	Valqueire News	3/16/15	9,680	358
204	Vaz Lobo News	2/8/16	21,501	909
205	Vicente de Carvalho em Depressão	3/14/13	15,346	244
206	Vidigal News	7/21/11	21,603	2,641
207	Vigírio Geral	5/14/19	3,279	172
208	Vila Autódromo	4/19/12	9,986	1,312
209	Vila da Penha News	9/3/15	5,389	6,872
210	Vila da Penha Notícias	8/26/17	57,835	2,960
211	Vila da Penha Online	2/17/11	10,146	214
212	Vila Kennedy RJ	10/3/11	43,119	1,816
213	Vila Kosmos	12/12/11	7,042	2,644
214	Viva Rocinha	6/4/11	23,086	1,720
215	VOZ da Vila Kennedy	6/11/17	87,074	4,270
216	Voz das Comunidades	4/14/11	199,213	13,799
217	Voz de Piedade	5/22/15	10,686	1,428
218	Voz de Quintino	6/12/15	36,373	7,032
219	Voz do Engenho de Dentro	4/10/15	37,215	3,196
220	Zona Portuária Alerta	9/4/13	20,537	6,325
221	Zona Sul da Depressão	2/20/14	42,030	2,201
	221 groups		**8,584,818 followers**	**610,389 posts**

Table 7.2: News Media Pages

#	Outlet Name	Followers (2020)	# Posts	Local?
1	Agora São Paulo	8,499	3,226	✓
2	Band FM	1,322,200	45,683	
3	BandNews TV	869,894	31,366	
4	CNNBrasil	723,651	98	
5	Diário do Rio de Janeiro	359,320	13,570	✓
6	Estadão	3,671,655	121,755	
7	Folha de S.Paulo	5,661,269	78,351	
8	G1 - O Portal de Notícias da Globo	10,804,667	73,466	
9	GloboNews	2,498,040	33,559	
10	Jornal da Record	674,332	15,351	
11	Jornal Extra	2,728,398	84,588	
12	Jornal O Dia	680,391	59,204	
13	O Globo	5,649,413	109,012	
14	O POVO Online	1,385,428	96,971	
15	Rede Globo	13,925,450	49,541	
16	RedeTV!	4,836,204	137,495	
17	Rádio BandNews FM - Rio de Janeiro	232,016	34,995	✓
18	Rádio Rio de Janeiro	27,842	8,708	✓
19	SBT	11,507,331	40,467	
20	Terra	5,039,482	112,275	
21	UOL	7,977,499	204,308	

Table 7.3: Violence Search Terms

	Terms
Original Terms	"tiro","armas", "bala","fogo", "pistola", "balead", "atingid", "crimino", "delito" "crime","violac","ofensa","violencia", "violent", "abuso", "desordem", "raiva" "tumulto", "alvoroco", "luta", "assalto", "ataque", "estupro", "extorsao", "roub" "incursao", "infracao", "rapto","morte", "falecimento", "decadencia","morre", "falec" "morto","mata", "homicíd", "assassinato", "assassino", "lesao", "ferid","ferimento" "latrocinio", "furto", "ladroes", "ladrao", "bandid", "sequestro", "abdutor", "desaparecid" "trafico de droga","narcotrafic", "traficante", "droga", "gangue", "comando vermelho" "crime organizado", "faccoes","milicia", "polici", "blitz", "operac", "upp", "pm" "batal", "bope", "pmerj", "seguranca","militariza"
Doc2Vec Terms	"estelionato","envolvimento","pedofilia","torturar","feminicidio" "acusado","chefiar", "quadrilha", "acusada","tortura", "suspeito", "confessa", "hediondos" "cometido", "importunacao", "agredir", "acusado", "flagrante", "suspeito", "prisao" "assassinato", "suspeita", "vitima", "estelionato","investigacoes", "acusada", "extorsao", "criminalidade", "caos", "medo", "sofrem", "criminalizacao", "mortes" "guerra", "armada"

References

Adena, Maja, Ruben Enikolopov, Maria Petrova, Veronica Santarosa and Ekaterina Zhuravskaya. 2015. "Radio and the Rise of the Nazis in Prewar Germany." *The Quarterly Journal of Economics* 130(4):1885–1939.

Agre, Philip E. 2002. "Real-time politics: The Internet and the political process." *The information society* 18(5):311–331.

Ajzen, Icek and Martin Fishbein. 1977. "Attitude-behavior relations: A theoretical analysis and review of empirical research." *Psychological bulletin* 84(5):888.

Allcott, Hunt, Luca Braghieri, Sarah Eichmeyer and Matthew Gentzkow. 2019. The welfare effects of social media. Technical report National Bureau of Economic Research.

Allcott, Hunt and Matthew Gentzkow. 2017. "Social media and fake news in the 2016 election." *Journal of economic perspectives* 31(2):211–36.

Allison, Paul D. 1990. "Change scores as dependent variables in regression analysis." *Sociological methodology* pp. 93–114.

Altheide, David L. 1997. "The news media, the problem frame, and the production of fear." *The sociological quarterly* 38(4):647–668.

Altheide, David L. 2002. *Creating fear: News and the construction of crisis*. Transaction Publishers.

Ansolabehere, Stephen and Shanto Iyengar. 1995. *Going negative*. Vol. 95 New York: Free Press.

Arias, Eric, Horacio Larreguy, John Marshall and Pablo Querubin. 2018. Priors Rule: When do Malfeasance Revelations Help or Hurt Incumbent Parties? Technical report National Bureau of Economic Research.

Asimovic, Nejla, Jonathan Nagler, Richard Bonneau and Joshua A Tucker. 2021. "Testing the effects of Facebook usage in an ethnically polarized setting." *Proceedings of the National Academy of Sciences* 118(25).

Bail, Christopher A, Lisa P Argyle, Taylor W Brown, John P Bumpus, Haohan Chen, MB Fallin Hunzaker, Jaemin Lee, Marcus Mann, Friedolin Merhout and Alexander Volfovsky. 2018. "Exposure to opposing views on social media can increase political polarization." *Proceedings of the National Academy of Sciences* 115(37):9216–9221.

Banerjee, Abhijit, Selvan Kumar, Rohini Pande and Felix Su. 2011. "Do Informed Voters Make Better Choices? Experimental Evidence from Urban India." *Unpublished manuscript* .

Bebbington, Keely, Colin MacLeod, T Mark Ellison and Nicolas Fay. 2017. "The sky is falling: evidence of a negativity bias in the social transmission of information." *Evolution and Human Behavior* 38(1):92–101.

Bechmann, Anja and Stine Lomborg. 2013. "Mapping actor roles in social media: Different perspectives on value creation in theories of user participation." *New media & society* 15(5):765–781.

Borges, Doriam. 2011. *O medo do crime na cidade do Rio de Janeiro: um análise sob a perspectiva das crençcas de perigo*. Editora Appris.

Bott, Maja, Björn-Sören Gigler and Gregor Young. 2014. The role of crowdsourcing for better governance in fragile state contexts.

Brannon, Laura A, Michael J Tagler and Alice H Eagly. 2007. "The moderating role of attitude strength in selective exposure to information." *Journal of Experimental Social Psychology* 43(4):611–617.

Breuer, Anita, Todd Landman and Dorothea Farquhar. 2015. "Social media and protest mobilization: Evidence from the Tunisian revolution." *Democratization* 22(4):764–792.

Brierley, Sarah, Eric Kramon and George Kwaku Ofosu. 2018. "The moderating effect of debates on political attitudes." *American Journal of Political Science* .

Buringh, Eltjo and Jan Luiten Van Zanden. 2009. "Charting the" Rise of the West": Manuscripts and Printed Books in Europe, A Long-Term Perspective from the Sixth through Eighteenth Centuries." *The Journal of Economic History* pp. 409–445.

Cagé, Julia. 2020. "Media competition, information provision and political participation: Evidence from French local newspapers and elections, 1944–2014." *Journal of Public Economics* 185:104077.

CESeC, Movimentos. 2019. "# FavelaLivesMatter: Youth from Urban Peripheries, Political Engagement and Alternatives to the War on Drugs." *Journal of Illicit Economies and Development* 1(2).

Chong, Alberto, Ana L De La O, Dean Karlan and Leonard Wantchekon. 2014. "Does Corruption Information Inspire the Fight or Quash the Hope? A Field Experiment in Mexico on Voter Turnout, choice, and party identification." *The Journal of Politics* 77(1):55–71.

Cobb, Michael D and James H Kuklinski. 1997. "Changing Minds: Political Arguments and Political Persuasion." *American Journal of Political Science* pp. 88–121.

Cohen, Bernard Cecil. 1963. *Press and foreign policy*. Princeton University Press.

Contrera, Malena Segura. 2002. *Mídia e pânico: saturação da informação, violência e crise cultural na mídia*. Annablume.

DellaVigna, Stefano and Ethan Kaplan. 2007. "The Fox News effect: Media bias and voting." *The Quarterly Journal of Economics* 122(3):1187–1234.

DiNucci, Darcy. 1999. "Fragmented future." *Print* 53(4):32–33.

Ditto, Peter H and David F Lopez. 1992. "Motivated skepticism: Use of differential decision criteria for preferred and nonpreferred conclusions." *Journal of Personality and Social Psychology* 63(4):568.

Dowd, Caitriona, Patricia Justino, Roudabeh Kishi and Gauthier Marchais. 2020. "Comparing 'New'and 'Old'Media for Violence Monitoring and Crisis Response: Evidence from Kenya." *Research & Politics* 7(3):2053168020937592.

Druckman, James N and Arthur Lupia. 2000. "Preference formation." *Annual Review of Political Science* 3(1):1–24.

Dunning, Thad, Guy Grossman, Macartan Humphreys, Susan D Hyde, Craig McIntosh, Gareth Nellis, Claire L Adida, Eric Arias, Clara Bicalho, Taylor C Boas et al. 2019. "Voter information campaigns and political accountability: Cumulative findings from a preregistered meta-analysis of coordinated trials." *Science Advances* 5(7):eaaw2612.

Eisenhardt, Kathleen M. 1989. "Agency theory: An assessment and review." *Academy of management review* 14(1):57–74.

Erete, Sheena L. 2015. Engaging around neighborhood issues: How online communication affects offline behavior. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing*. ACM pp. 1590–1601.

Farrell, Henry. 2012. "The consequences of the internet for politics." *Annual review of political science* 15.

Feltwell, Tom, John Vines, Karen Salt, Mark Blythe, Ben Kirman, Julie Barnett, Phillip Brooker and Shaun Lawson. 2017. "Counter-discourse activism on social media: The case of challenging "poverty porn" television." *Computer Supported Cooperative Work (CSCW)* 26(3):345–385.

Ferwerda, Jeremy, DJ Flynn and Yusaku Horiuchi. 2017. "Explaining Opposition to Refugee Resettlement: The Role of NIMBYism and Perceived Threats." *Science advances* 3(9):e1700812.

Festinger, Leon. 1957. *A Theory of Cognitive Dissonance*. Stanford University Press.

Fuchs, Christian. 2017. *Social media: A critical introduction*. Sage.

Funk, Jeanne B, Heidi Bechtoldt Baldacci, Tracie Pasold and Jennifer Baumgardner. 2004. "Violence exposure in real-life, video games, television, movies, and the internet: is there desensitization?" *Journal of adolescence* 27(1):23–39.

Gaines, Brian J, James H Kuklinski, Paul J Quirk, Buddy Peyton and Jay Verkuilen. 2007. "Same facts, different interpretations: Partisan motivation and opinion on Iraq." *Journal of Politics* 69(4):957–974.

Gao, Huiji, Geoffrey Barbier and Rebecca Goolsby. 2011. "Harnessing the crowdsourcing power of social media for disaster relief." *IEEE Intelligent Systems* 26(3):10–14.

Gerber, Alan and Donald Green. 1999. "Misperceptions about perceptual bias." *Annual review of political science* 2(1):189–210.

Gerber, Alan S, Dean Karlan and Daniel Bergan. 2009. "Does the media matter? A field experiment measuring the effect of newspapers on voting behavior and political opinions." *American Economic Journal: Applied Economics* 1(2):35–52.

Gerber, Alan S. and Donald P. Green. 2012. *Field Experiments: Design, Analysis, and Interpretation*. W.W. Norton.

Gillespie, Tarleton. 2010. "The politics of 'platforms'." *New media & society* 12(3):347–364.

Gleick, James. 2011. *The information: A history, a theory, a flood*. Vintage.

Grupe, Dan W and Jack B Nitschke. 2013. "Uncertainty and anticipation in anxiety: an integrated neurobiological and psychological perspective." *Nature Reviews Neuroscience* 14(7):488–501.

Guess, Andrew, Jonathan Nagler and Joshua Tucker. 2019. "Less than you think: Prevalence and predictors of fake news dissemination on Facebook." *Science advances* 5(1):eaau4586.

Guriev, Sergei, Nikita Melnikov and Ekaterina Zhuravskaya. 2019. "3g internet and confidence in government." *Available at SSRN 3456747* .

Hainmueller, Jens, Daniel J Hopkins and Teppei Yamamoto. 2014. "Causal inference in conjoint analysis: Understanding multidimensional choices via stated preference experiments." *Political analysis* 22(1):1–30.

Harari, Yuval Noah. 2014. *Sapiens: A brief history of humankind*. Random House.

Howe, Jeff. 2006. "The Rise of Crowdsourcing." *Wired* .
URL: *https://www.wired.com/2006/06/crowds/*

Humphreys, Macartan and Jeremy Weinstein. 2012. "Policing politicians: citizen empowerment and political accountability in Uganda preliminary analysis." *Columbia Universities. Unpublished manuscript* .

Israni, Aarti, Sheena Erete and Che L Smith. 2017. Snitches, Trolls, and Social Norms: Unpacking Perceptions of Social Media Use for Crime Prevention. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing*. ACM pp. 1193–1209.

Ito, Tiffany A, Jeff T Larsen, N Kyle Smith and John T Cacioppo. 1998. "Negative information weighs more heavily on the brain: the negativity bias in evaluative categorizations." *Journal of personality and social psychology* 75(4):887.

Iyengar, Shanto and Donald R Kinder. 2010. *News that matters: Television and American opinion*. University of Chicago Press.

Jensen, Klaus Bruhn. 2015. "What's social about social media?" *Social Media + Society* 1(1):1–2.

Jungherr, Andreas, Gonzalo Rivero and Daniel Gayo-Avello. 2020. *Retooling politics: How digital media are shaping democracy*. Cambridge University Press.

Kahneman, Daniel and Amos Tversky. 1979. Prospect theory: An analysis of decision under risk. In *Handbook of the fundamentals of financial decision making: Part I*. World Scientific pp. 99–127.

Kaplan, Andreas M and Michael Haenlein. 2010. "Users of the world, unite! The challenges and opportunities of Social Media." *Business horizons* 53(1):59–68.

King, Gary, Jennifer Pan and Margaret E Roberts. 2013. "How Censorship in China Allows Government Criticism but Silences Collective Expression." *American Political Science Review* 107(2):326–343.

Krosnick, Jon A and Robert P Abelson. 1992. "The case for measuring attitude strength in surveys." *Questions about questions: Inquiries into the cognitive bases of surveys* pp. 177–203.

Kuth, Andreas. 2006. "Among the audience." *The Economist* .
URL: *https://www.economist.com/special-report/2006/04/22/among-the-audience*

Lee, Jae Kook, Jihyang Choi, Cheonsoo Kim and Yonghwan Kim. 2014. "Social media,

network heterogeneity, and opinion polarization." *Journal of communication* 64(4):702–722.

Levine, Lawrence W and Cornelia R Levine. 2010. *The fireside conversations: America responds to FDR during the Great Depression*. Univ of California Press.

Lieberman, Evan S, Daniel N Posner and Lily L Tsai. 2014. "Does Information Lead to More Active Citizenship? Evidence from an Education Intervention in Rural Kenya." *World Development* 60:69–83.

Lupia, Arthur. 1994. "Shortcuts versus encyclopedias: Information and voting behavior in California insurance reform elections." *American Political Science Review* 88(1):63–76.

Mäkinen, Maarit and Mary Wangu Kuira. 2008. "Social media and postelection crisis in Kenya." *The International Journal of Press/Politics* 13(3):328–335.

Margetts, Helen, Peter John, Scott Hale and Taha Yasseri. 2015. *Political turbulence: How social media shape collective action*. Princeton University Press.

McGuire, William J. 1985. "Attitudes and attitude change." *The handbook of social psychology* pp. 233–346.

McLaughlin, Katie A and Mark L Hatzenbuehler. 2009. "Stressful life events, anxiety sensitivity, and internalizing symptoms in adolescents." *Journal of abnormal psychology* 118(3):659.

Meier, Patrick. 2011. "Theorizing Ushahidi: An Academic Treatise." *posted on the iRevolution blog* .

Moors, M Rae. 2019. "What is Flint? Place, storytelling, and social media narrative reclamation during the Flint water crisis." *Information, Communication & Society* 22(6):808–822.

Nahon, Karine, Jeff Hemsley, Shawn Walker and Muzammil Hussain. 2011. "Fifteen minutes of fame: The power of blogs in the lifecycle of viral political information." *Policy & Internet* 3(1):1–28.

Nunes, Pablo. 2017. "Crime e Polícia no RiodeJaneiro: Relatos Em Páginas do Facebook." *Boletim Segurança e Cidadania* .

Nye, Joseph S, Andrei Cherny and Eric Alterman. 1990. *What liberal media?: The truth about bias and the news*. Vol. 10 Basic Books.

OECD. 2009. *Preventing violence, war and state collapse: The future of conflict early warning and response*. Organisation for Economic Co-operation and Development.

O'reilly, Tim. 2005. "What is web 2.0.".

Popkin, Samuel L and Samuel L Popkin. 1994. *The reasoning voter: Communication and persuasion in presidential campaigns*. University of Chicago Press.

Poushter, Jacob, Caldwell Bishop and Hanyu Chwe. 2018. "Social media use continues to rise in developing countries but plateaus across developed ones." *Pew research center* 22:2–19.

Pridmore, Jason, Anouk Mols, Yijing Wang and Frank Holleman. 2019. "Keeping an eye on the neighbours: Police, citizens, and communication within mobile neighbourhood crime prevention groups." *The Police Journal* 92(2):97–120.

Ramos, Silvia and Anabela Paiva. 2007. *Mídia e violência*. Iuperj Rio de Janeiro.

Ramos, Silvia, Anabela Paiva and Pablo Nunes. 2013. *Mídia e violência: O que Mudou em uma década*. Centro de Estudos de Segurança e Cidadania - CESeC.

URL: https://www.ucamcesec.com.br/wp-content/uploads/2016/10/Midia-e-Violencia_2016_Versao4.pdf

Redlawsk, David P, Andrew JW Civettini and Karen M Emmerson. 2010. "The affective tipping point: Do motivated reasoners ever "get it"?" *Political Psychology* 31(4):563–593.

Rentschler, Eric. 2003. "The Fascination of a Fake: The Hitler Diaries." *New German Critique* (90):177–192.

Rigterink, Anouk S and Mareike Schomerus. 2017. "The fear factor is a main thing: How radio influences anxiety and political attitudes." *The Journal of Development Studies* 53(8):1123–1146.

Roberts, Tony and Gauthier Marchais. 2017. *Assessing the Role of Social Media and Digital Technology in Violence Reporting*. IDS.

Rosen, Jay. 2008. The people formerly known as the audience. In *Participation and media production: Critical reflections on content creation*. Cambridge Scientific Publishers Ltd.

Rozin, Paul and Edward B Royzman. 2001. "Negativity bias, negativity dominance, and contagion." *Personality and social psychology review* 5(4):296–320.

Shapiro, Jacob N and Nils B Weidmann. 2015. "Is the Phone Mightier than the Sword? Cellphones and Insurgent Violence in Iraq." *International Organization* 69(2):247–274.

Shirky, Clay. 2008. *Here comes everybody: The power of organizing without organizations*. Penguin.

Singer, Peter Warren and Emerson T Brooking. 2018. *LikeWar: The Weaponization of Social Media*. Eamon Dolan Books.

Sniderman, Paul M, Richard A Brody and Phillip E Tetlock. 1993. *Reasoning and choice: Explorations in political psychology*. Cambridge University Press.

Sobolev, Anton. 2018. "How Pro-Government "Trolls" Influence Online Conversations in Russia." *University of California Los Angeles.[7]* .

Soroka, Stuart N. 2006. "Good news and bad news: Asymmetric responses to economic information." *The journal of Politics* 68(2):372–385.

Surowiecki, James. 2004. "The wisdom of crowds: Why the many are smarter than the few and how collective wisdom shapes business." *Economies, Societies and Nations* 296:5.

Taber, Charles S and Milton Lodge. 2006. "Motivated skepticism in the evaluation of political beliefs." *American Journal of Political Science* 50(3):755–769.

Talbot, Adam. 2019. "Talking about the 'rotten fruits' of Rio 2016: Framing mega-event legacies." *International Review for the Sociology of Sport* p. 1012690219878842.

Tormala, Zakary L and Derek D Rucker. 2007. "Attitude certainty: A Review of Past Findings and Emerging Perspectives." *Social and Personality Psychology Compass* 1(1):469–492.

Trussler, Marc and Stuart Soroka. 2014. "Consumer demand for cynical and negative news frames." *The International Journal of Press/Politics* 19(3):360–379.

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):363–379.

Tversky, Amos and Daniel Kahneman. 1974. "Judgment under uncertainty: Heuristics and biases." *science* 185(4157):1124–1131.

Unver, A and H Alasaad. 2016. "How Turks Mobilized against the Coup: The Power of the Hashtag and the Mosque." *Foreign Affairs* 14.

Van der Windt, Peter and Macartan Humphreys. 2016. "Crowdseeding in Eastern Congo: Using cell phones to collect conflict events data in real time." *Journal of Conflict Resolution* 60(4):748–781.

Walsh, James P. 2020. "Social media and moral panics: Assessing the effects of technological change on societal reaction." *International Journal of Cultural Studies* .

Wolfsfeld, Gadi, Elad Segev and Tamir Sheafer. 2013. "Social media and the Arab Spring: Politics comes first." *The International Journal of Press/Politics* 18(2):115–137.

Zaller, John and Dennis Chiu. 1996. "Government's little helper: US press coverage of foreign policy crises, 1945–1991." *Political Communication* 13(4):385–405.

Zaller, John R. 1992. *The nature and origins of mass opinion*. Cambridge university press.

Zeitzoff, Thomas. 2017. "How Social Media is Changing Conflict." *Journal of Conflict Resolution* 61(9):1970–1991.

Zhuravskaya, Ekaterina, Maria Petrova and Ruben Enikolopov. 2020. "Political effects of the internet and social media." *Annual Review of Economics* 12.