The Cop In Your Neighbor's Doorbell

Amazon Ring And The Spread of Participatory Mass Surveillance

ANONYMOUS AUTHOR(S)

Consumer surveillance products such as 'smart' doorbell cameras are an already-pervasive phenomenon in the U.S. These devices are marketed as personal and community security tools that allow users to answer their front door remotely, record "suspicious activity" captured by their cameras, and share reports with neighbors. The widespread use of doorbell cameras specifically, however, has created an opaque, wide-reaching surveillance network used by thousands of law enforcement agencies nationwide [11, 19, 44]. The full breadth of this network and how users operate on such platforms is largely unknown. Amazon Ring, one of the largest manufacturers of smart doorbells, offers a companion social networking app to their physical doorbells called Ring Neighbors that allows camera owners to share video and text posts with other camera owners that live nearby. In this paper, we use data collected from public posts on Neighbors to create what we believe is the first comprehensive map and analysis of smart doorbell camera use across the continental U.S. We use spatial regression methods to estimate the county-level predictors of Neighbors app usage nationally. We then use Los Angeles, one of the most active areas of Ring usage in the country, as a case study to investigate how different neighborhoods in a racially heterogeneous city use a platform like Ring. Using a structured topic analysis and experimental survey design, we show that users actively frame video subjects as criminal and suspicious, that the race of a neighborhood has a significant impact on posting rates, and some evidence that Neighbors may be used as a racial gatekeeping tool, particularly by white neighborhoods that border non-white areas in Los Angeles.

CCS Concepts: • **Social and professional topics** → **Corporate surveillance**; • **Applied computing** → *Sociology*.

Additional Key Words and Phrases: Data & Society, Surveillance, platforms

ACM Reference Format:
Anonymous Author(s). 2021. The Cop In Your Neighbor's Doorbell: Amazon Ring And The Spread of Participatory Mass Surveillance. 1, 1 (April 2021), 56 pages.

ACKNOWLEDGEMENTS

1 INTRODUCTION

Online social platforms are ubiquitous in American life, with participation online increasingly becoming a core part of social and civic engagement [16, 32]. As online platforms grow as hubs of neighborhood and civic activity, researchers have studied the potential harms of platform design, including polarization, hate speech, disinformation, and extremism [49]. Although these topics have been studied in depth in online communities on platforms like Facebook and Twitter, less is known about the dynamics of "hyper-local" social networks that are designed as digital analogues to neighborhood connection and civic engagement, such as Nextdoor and Ring Neighbors [23, 35]. But much as online platform design has been shown to impact the quality of civic and media engagement on sites like Facebook, these hyper-local platforms carry with them a risk of

© 2022 Association for Computing Machinery.
Manuscript submitted to ACM

negatively impacting community dynamics or reinforcing existing harmful community behavior, such as race-based community gatekeeping [23].

To add to this discourse, we study Amazon's Ring Neighbors platform, a hyper-local social networking app that is built around the spread and use of 'smart' security and video surveillance products offered by Ring. Their most common product, the Ring Doorbell, offers owners live and motion-activated surveillance footage of their front doorsteps, and their app encourages owners to share recorded videos with their community through the Ring Neighbors app. Videos shared on the platform or recorded by devices are not just used by other camera owners, however. Ring partners with over 2,000 law enforcement agencies nationwide, giving them access to community networks and the ability to obtain videos from users [45]. These unique features put into sharp focus the risks of integrating online platforms into real-world community governance—on Ring, sharing a video of a stranger on your doorstep can spur police action. This high level of potential impact has brought Ring under intense public scrutiny, with journalistic investigations finding that Ring video posts disproportionately portray people of color as suspicious [19], that they influence people's perceptions of their neighborhoods [29], and that Ring uses their police networks to sell products [11].

In this paper, we aim to characterize the socio-demographic factors that make a neighborhood more likely to use the Ring Neighbors platform and investigate how neighborhood demographics impact the kinds of alerts people post. To get a bird's-eye view of how people across the U.S. use the platform, we use a mixed methods approach. First, we present a critical summary of Ring as a platform, and build on surveillance studies literature to frame the technology within long-standing conversations on surveillance and community self-policing. We then shift to a data-driven approach, and use data scraped from the app between October 2016 and April 2020 to map the location of all the users that have published posts in that nearly four-year period. We then use this geotagged data to estimate the impact of county socio-demographic attributes on heightened rates of platform use, finding patterns that are at odds with some of the ways the Ring technology is often marketed and described.

We use the city of Los Angeles as a case study in how a specific urban area uses the platform, and present a spatial regression analysis testing hypotheses that reflect a common journalistic analysis: that Ring enables and encourages racial and economic gatekeeping in neighborhoods. After finding significant racial and economic impacts on Ring usage, we leverage computationally grounded methods [33], including structural topic models, to investigate how demographics mediate different uses of the platform.

1.1 Ring Neighbors and the Platform

Neighbors is a social media platform accompanying Amazon's ecosystem of cameras, flood lights, and other Internet of Things devices. According to Ring's website, the app allows people to "connect with [their] neighbors and stay up-to-date with what's going on in [their] neighborhood" [4]. Framed as "the new neighborhood watch", the app contains a social feed and map of posts uploaded in the general vicinity (up to 8km away) by users, local law enforcement, or the Ring moderators (see Figure 1). Unlike other hyper-local neighborhood social networks such as Nextdoor, posts on Ring Neighbors are primarily related to crime and public safety—users can select one of six categories (safety, crime, lost pet, unexpected activity, neighborly moment, or "I'm not sure") when uploading posts. Additionally, all posts on Ring are moderated [43]. Each post also contains a title, description, up to five photos or videos, and a location, anonymized to a nearby street intersection. Like other social networks, users can upvote, share, or comment on alerts,

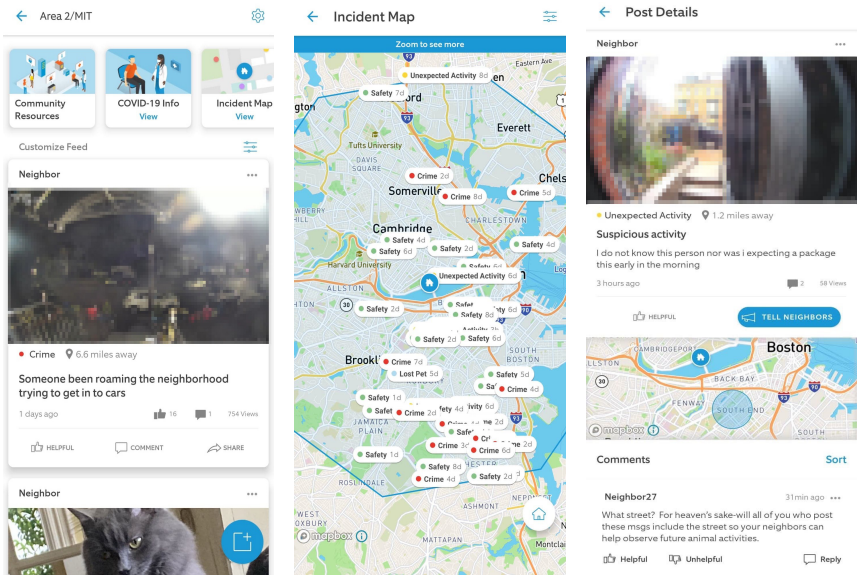

Fig. 1. Left: A feed of posts on the Neighbors app. Center: Neighbors posts shown on the app's map. Right: An example post on the Neighbors app. Screenshot taken by the authors.

except all user names are anonymized (see Figure 1). Users can also customize the size and shape of their "neighborhood"—the geographic region they see posts from—as well as sign up to receive real-time push alerts for new posts.

Ring, and the Neighbors app in particular, has drawn varied attention and critique in recent years. Although Ring's terms of service dictates that customers install their cameras in a way that records only their personal property—i.e. their front porch—this is largely ignored by its users. Most cameras have a clear view into neighoring property, streets, sidewalks, parks in front of the home, and more. As journalist Caroline Haskins points out, when a consumer purchases and installs a Ring doorbell, they "make a decision on behalf of everyone around [them]. If someone walks by [their] house, lives next door, or delivers packages to [their] home, they will be recorded and surveilled. They don't get a choice" [20]. Several journalistic investigations have found that posts on the Neighbors platform disproportionately depict people of color [12, 19], raising concerns that the platform furthers inequality and racism in policing practices through "constructing a web of police surveillance" [17] whose gaze is trained primarily on people of color. We argue that the impact of an individual consumer that participates on the Neighbors platform extends beyond even the neighbors and visitors recorded by the platform. It places the Ring user in the role of a "prosumer" [40] of surveillant content, furthers the responsibilization of citizens as agents of law enforcement, reproduces patterns of racial gatekeeping in neighborhoods around the U.S., and advances a new form of surveillant relationship: participatory mass surveillance.

1.2 Ring Neighbors as Prosumer Policing

In the context of Ring as a social media platform, users of Ring can be understood as what George Ritzer and Nathan Jurgenson refer to as *prosumers*, where the users of a product are both consumers and producers simultaneously [40]. On the Ring Neighbors app, users produce and consume content such as text, images, and videos of "suspicious" Others, lost pets, and crime. In this way, Ring continues a trend of online platforms "putting consumers to work", but instead of producing entertainment content as in other platforms, prosumers on Ring produce *surveillant content* [48]. Like other prosumer platforms such as Facebook or Twitter, the primary consumers of the content produced are other users of the platform. However, in many other platform arrangements, the secondary consumer of user content is the platform itself: data on consumer behavior is consumed by platforms to inform algorithm or process development to improve a product or influence consumer behavior [51]. In the case of Ring, however, users create content with the implicit knowledge that it will be consumed—and potentially acted on—by law enforcement.

This contextualizes Ring as a special kind of prosumer arrangement consistent with a modern trend of *responsibilisation* in modern policing. Responsibilisation refers to the phenomenon of citizens tasked, informally or formally, with performing certain policing or carceral logics themselves, such as surveillance and reporting [31, 46, 50]. The surveillant content produced by users can be seen as a way users adopt some policing responsibility: Ring describes posting videos of people and "suspicious" activity as users "doing their duty" and frames participating in their law enforcement partnerships, where users can share videos directly with police, as providing a valuable service to law enforcement [11]. By facilitating the responsibilisation of policing practices, Ring operates as a kind of carceral technology that operates at the level of the neighborhood.

1.3 "The New Neighborhood Watch" and Digital Community Gatekeeping

As part of this responsibilisation, Ring also describes itself as "the new neighborhood watch". This references the long history of neighborhood watch movements in the U.S. and beyond: community-organized policing efforts that enforce local norms and laws that operate outside of and extend formal policing structures [3, 9, 17]. Maya Schenwar and Victoria Law argue that civilian policing in the U.S. originated as a way to maintain white supremacy: for instance, many southern police departments started as community organized slave patrols, and volunteer police helpers facilitated Indigenous genocide. In the 1960s, the neighborhood watch, rooted in these civilian policing movements, gained popularity and encouraged residents to report suspicious behavior to the police [24]. While it is unlikely Ring employs this language to specifically call to mind this historical context as a dog whistle, their use of this shared language is notable. Modern neighborhood watches facilitated by other technologies like NextDoor, another 'hyperlocal' neighborhood app, and communication platforms like WhatsApp, have been found to engage in racial or ethnic profiling and vigilantistic policing practices. [23, 31].

Rahim Kurwa argues that neighborhood apps like NextDoor operate under "quasi-carceral" logics that apply a criminalizing gaze to Black and other non-white community members, enabling white residents to "digitally gate" neighborhoods through private policing [23]. The platform we study here, Ring Neighbors, is similar to NextDoor, and also encourages these quasi-carceral logics of community policing, but is more explicitly connected to formal modes of policing via law enforcement partnerships. As such, Ring can be

considered a carceral technology that enforces racial boundaries through both informal "neighborhood watch"-style mechanisms and formal policing. Both of these avenues can be ways for racially homogeneous neighborhoods to enforce segregatory preferences and dictate who belongs in their community—what we call "racial gatekeeping". We investigate this possible use of Ring in our analysis and reading of Neighbors posts in Los Angeles.

1.4 Social and Participatory Mass Surveillance

Ring's prosumer and policing contexts also make its model of surveillance unique. Unlike CCTV, a common object of past surveillance literature which Kurwa points out is "depersonalized, unidirectional in its gaze from camera to subjects, and needs to be read and interpreted by a human being", surveillance shared on the Ring platform is personal, and any content shared is interpreted by the owner of the camera immediately through text that accompanies posts on the platform [23]. The personalized nature of Ring's surveillance suggests that it might be categorized as lateral or dyadic surveillance: the target of surveillance is anyone wandering into the Ring camera's frame, and the "watcher" is the camera owner [7]. But this relationship does not fully characterize relations within the platform. Surveillance and privacy are best described in relations of power, not simply patterns of observation [27].

When a video of someone is shared by a user on the Ring platform, that user also implicitly enters the subject into a surveillant power relationship with the owner, other users on the platform, Amazon, and with local police all at once. This gives device owners considerable power over those being recorded—they are the ones who ultimately decide which videos are uploaded to the platform, and which videos, through Ring's law enforcement partnerships, are made available to police. In this way, users of Ring Neighbors, by posting on the app or by simply recording video that may later be given to law enforcement, expand the circle of control that local police have over their community.

At the same time, as Lauren Bridges argues, Ring should be situated within the context of Amazon and Amazon Web Services, a massive digital infrastructure that has historically provided facial recognition software to police, and is known for leveraging data-mining techniques to extract value from user data [9, 51]. In 2018, *The Washington Post* reported on a patent filed by Amazon that showed designs for a massive database of "suspicious persons", automatically identified through facial recognition software applied to Ring's video surveillance [21]. The capacity for automatically data-mining videos collected through Ring using facial recognition or other tools offers another surveillant model that needs to be considered: one of participatory mass surveillance, where individual lateral surveillance practices not only emulate and amplify state surveillance practices, but also contribute to a unidirectional surveillance apparatus controlled by a private corporation and shared with the state [7].

This potential extension has worried privacy advocates and legal experts as a potential erosion of Fourth Amendment rights [47]. As legal scholar Joel Reidenberg argues, modern interpretations of a 'reasonable expectation to privacy' rest on what technology—and what information—is currently commonly commercially deployed or available to consumers. The future of Ring may mean that everyday consumers have access to technology that transmits recordings and biometric data of people on porches, streets, and sidewalks to a central database, all made available to police. If Ring continues to grow and these concerns go unanswered, where might Fourth Amendment rights begin and end? [37].

2 EMPIRICAL QUESTIONS AND RING'S NATIONAL REACH

2.1 Data Collection and Description

We have argued that surveillance and policing theory can frame Ring as an extension of policing into private communities, as a continuation of "neighborhood watch" groups, and as a potential mechanism for neighborhood gatekeeping. While these theories and frames are useful ways to understand and contextualize Ring and the Neighbors app, they do not provide clarity into how Ring and the Neighbors platform are actually used. To test some of these theories and to provide an empirical grounding for the spread of the Ring network, we leverage a dataset of over 850,000 Ring alerts (posts) posted to the Neighbors app by over 650,000 unique users between October 2016 and February 2020 in the continental United States.

To collect the data, the authors developed an automated script that scraped posts on the platform by impersonating a user. The authors used only one user account to perform the scraping (so there was no impersonating of accounts). At the time of data collection, the Ring Neighbors app presented posts as an "infinite scroll" within a user's defined home location (see Figure 1), and API calls from the app were sent to a Ring API server unencrypted. To scrape post data, the script methodically changed the "home" region of the user and "scrolled" through all available posts through the API. Home regions were computed so as to entirely cover the continental United States. The resulting date range of posts ranges from October 2016 to February 2020 because the earliest post on the platform appears in October 2016, and scraping was performed in March 2020.

Each alert includes a unique ID, a user ID of the account the alert was posted by, a user-provided title, description, and category, as well as a timestamp and geocoded location. Example records from our dataset are presented in Appendix A. The locations are semi-anonymized by Ring, attributed to the nearest intersection instead of the exact location or address of the user or their device.

While this dataset is, to our knowledge, the first catalogue of Amazon Ring's surveillance network and product adoption that exists, we are careful not to claim this dataset as a representative sample of Ring camera owners or even Neighbors users. Because our data comes from public posts from the Ring Neighbors app, the dataset only includes data from users that have posted publicly using the platform; there are, by many counts, millions more Ring devices and Neighbors users than those devices and users that post to Ring Neighbors.

2.2 Methods

To characterize the national usage of Neighbors, we aggregate posts to the county and state level, and characterize each county by the number of alerts posted during our observation period per 1,000 residents as of the 2015-2018 ACS. Figure 2a shows this metric as recorded in each county in the continental United States. From observing the map, it is immediately clear both that Ring has national reach, with at least one alert posted in 64% of counties in the U.S., and that Ring Neighbors seems to be used more heavily in urban areas. Figure 2b shows the number of Ring alerts posted per day. At that time, the number of alerts per day on the platform (in all areas) was just below 500. By early 2020, that number had almost quadrupled, peaking at over 2,000 posts a day. Figure 2c shows a cumulative plot showing the number of states that reached 1,000 posts in total over time. By mid-2018, a year and a half after the first post we recorded (in California, October 2016), half of all states had reached at least 1,000 posts.

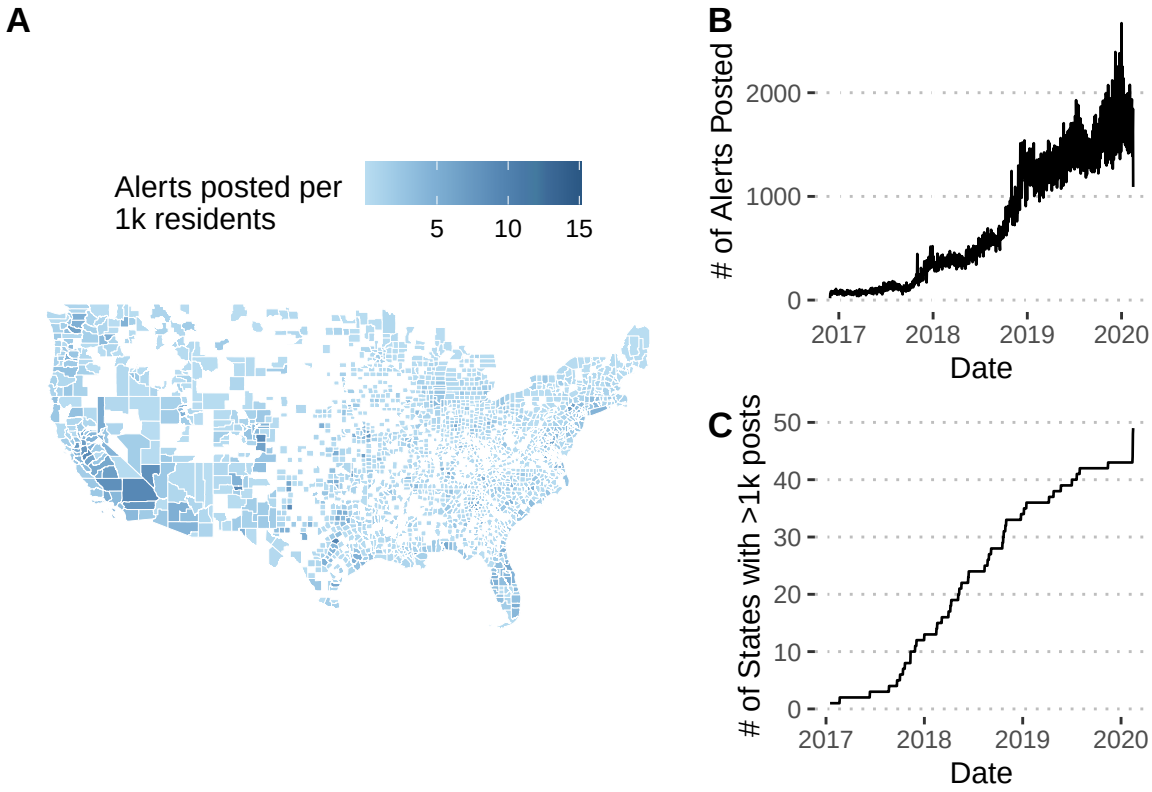

Fig. 2. The national reach of Amazon Ring. Figure A shows the number of alerts posted per-capita for each county in the continental United States. Counties that appear as white have no alerts. Figure B shows the number of posts on Ring Neighbors by day from October 2016 through February 2020, with a peak of 2671 alerts on January 01, 2020

To answer exploratory questions about the relationship of Ring Neighbors usage with area crime reporting and certain demographic variables, we use a spatial regression approach, creating independent variables that represent major categories of crime reporting and demographic attributes. Our crime reporting data is collected from the FBI UCR, a database of crime reporting statistics at the county level. We use crime reporting from 2017, as this is the first full year that Ring was used nationally [22].

For demographic variables, we use the 2015-2018 ACS, and extract variables related to property ownership and values, racial makeup, and income. For each county, we regress on its racial composition, median household income, homeownership rate, median home value, and crime reporting rates, including motor vehicle theft, property theft, robbery, assault, and total violent crime. We also include fixed effects for each state, and controls for the area of each county. Our outcome of interest, alerts per capita, effectively controls for total population. All our variables are mean-subtracted (centered) and divided by their standard deviation (standardized). We standardize instead of only center our data because the relative units of demographic percentages vary widely. For example, the standard deviation of the percentage of non-Hispanic white population in U.S. counties is 19.10%, while a standard deviation for the Native Hawaiian and Pacific Islander (NHPI) population is only 0.16%.

2.2.1 Spatial Modeling. Our covariates and our outcome variable have significant spatial dependence, which is not captured by a simple OLS regression model. To test dependence, we first create a network of counties by constructing a binary edge between two counties if they share a boundary, and then calculate Moran's I for the resulting OLS model. We calculate that our fully-specified model with all covariates has a Moran's I of 0.187 with p=1.1e-41, showing significant spatial dependence: unsurprisingly, the usage of Neighbors in one county is dependent on Neighbors usage in neighboring counties. To try and correct for this, we test the relative fit of different spatial regression models using an LM test, which shows that a spatial autoregressive lag model (as opposed to an error model) fits our data best (test statistic of 134.1 vs. 144.2). This makes sense: spatial lag models include spatial interactions as a set of additional terms, with the assumption that counties that are closer together have a greater impact on each other. An error model, on the other hand, treats any spatial dependence as an error to be modeled.

A spatial lag model's estimates are different than a traditional regression model. Instead of individual estimates for each covariate, a spatial lag model reports the overall impact that a variable has on an outcome of interest. Once a model is fit to the data, these impacts are estimated by measuring the effect that simulated changes on a covariate has on an outcome for both one observation (in our case, an individual county) and neighboring areas. For example, to estimate the impact of property ownership on the number of alerts per capita, we simulate an increase in property ownership in a county in the middle of Texas. Because spatial dependence is inter-linked, while the model will likely report a predicted change in our outcome variable in immediately neighboring counties, those counties' changes will have a ripple-effect-like impact on their neighboring counties, and so on. The impact for a variable is then estimated by performing this simulation using Markov Chain Monte Carlo draws for all variables for all counties several hundred times. This method also allows us to estimate the indirect and direct impact of a variable change on our outcome, measured as the average impact a change in one county may have on its neighbors, and the impact a change has on one county itself, respectively. In all results below, we report the estimated total impact of each variable, which is its indirect and direct impacts combined.

2.3 National Findings

There are a few clear hypotheses that may follow from our theoretical analysis of Ring as a platform. First, if Ring is used as a racial community gatekeeping tool, we may expect usage (alerts posted per-capita) to be overall higher in whiter counties, controlling for income and variables related to home ownership and property value. This simple test doesn't account for more complex spatial socio-dynamics, such as highly racially mixed counties that are very segregated, but it would indicate that there is a racial dynamic at play. We investigate the racial gatekeeping hypothesis more thoroughly in our quantitative case study of Los Angeles in Section 3.

Second, we expect property ownership rates, household median income, and property value to all have a positive impact on a county's usage of Neighbors. While Ring is much more affordable than traditional home security systems, purchasing and using one comes at a substantial cost. The devices themselves cost between $59.99 and $249.99, and a monthly subscription to store recorded video costs an additional $3-10 per month. Many of their products are more readily available to property owners, not renters, because they require altering property to install. They are also clearly marketed to people living in single-family, detached units.

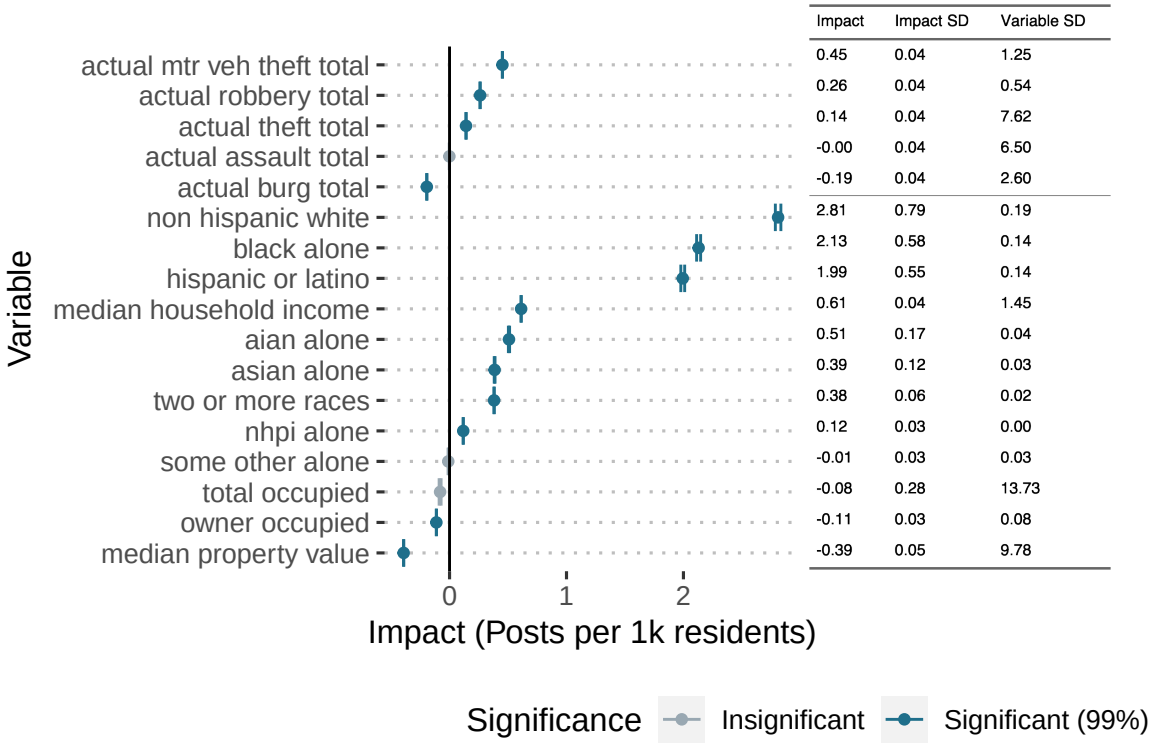

Fig. 3. Spatial Autoregressive Lag model results at the county level, including fixed effects by state. All variables are centered and standardized before modeling. Outcome variable is the number of Ring Neighbors posts per-capita in a county posted from October 2016 through February 2020 (mean 1.12, sd 1.5). Median household income and median property values are in tens of thousands of dollars. Total number of housing units occupied is in tens of thousands of units. All race-related variables are expressed as percentages of total population in a county. Crime reporting variables are reported as incidences per 1,000 people. Because our variables have been centered and standardized, the units for both our covariates and outcome variable are presented in standard deviations from their mean. See Appendix B for a description of each variable.

Third, we hypothesize that counties which report higher levels of property crime post on the Neighbors app more (controlling for the racial makeup of a county). Ring appears to be primarily a tool for policing and controlling property related crimes.

The results of this spatial modeling are visualized in Figure 3, with one row for each variable of interest. The fixed effect of a county's state, and its total population, are omitted for convenience. The impact estimates are expressed in units of standard deviations for each variable, including our outcome. For example, our results show that one standard deviation increase in the share of non-Hispanic white population in a county is associated with a 2.81 SD increase in the number of alerts per-capita. The SD of each variable in its original units is shown in the third column of the table. In the case of the non-Hispanic white share for a county, an increase of one SD is equivalent to 19.10%.

One of the most immediate observations from the spatial modeling results is the relative impact of the whiteness of a county on Neighbors use. A unit increase in the non-Hispanic white share in a county is

associated with an over 31% increase in posting rates compared to the next-highest racial variable, the share of Black residents in a county. This corresponds to an increase of 2.99 alerts posted per 1,000 people. This result provides evidence that usage of Ring is correlated with race—in particular, white communities use the platform more.

Interestingly, median property value and ownership rates have a *negative* impact on a county's posting rate, even while the median household income shows a strong positive impact. Higher income households are more likely to be able to afford nonessential consumer products like Ring, so this relationship is expected. We interpret the negative impact of the median property value in a county to suggest that households in counties with very high property values often already have other home security products, and so might not engage with the Ring ecosystem. The slight negative impact of ownership rates on posting rates contradicts our hypothesis that Ring products are more readily available to property owners than renters, though our scale (county) could be masking more complicated factors

Looking at the crime reporting variables included in our analysis, we see that all variables except for assault reporting rates are significant. While burglary rates have a negative impact, robbery rates have a positive impact of a similar magnitude. Burglary is defined as unlawful entry to commit a theft, while robbery strictly refers to theft under the threat of violence [2]. Previous research has shown that perceived rates of robbery are more likely than burglary to "generate fear" in communities, perhaps underlying a motivation for participating in the Neighbors network [28]. Also interesting is the relatively large impact of motor vehicle theft on posting rates. Motor vehicle theft is one of the least-common forms of property crime, occurring less than half as often as burglary or larceny, yet it is more strongly associated with Ring Neighbors posting rates than burglary or robbery combined. Content and anxieties about vehicle break-ins and theft are common on the platform (we discuss this in more detail in section 3 in our case study of Los Angeles), suggesting that perceived rates of vehicle-based crimes may be an important factor in community adoption of the platform. These results suggest that while property-based crime is a major factor in Neighbors usage, fear of personal violence in the form of robbery may also be a motivating factor.

This bird's eye analysis of the nation's use of Neighbors offers some important first clues as to why and how different communities use the Neighbors platform. Race plays a clear role in posting rates, with whiter communities posting more on the platform, and Neighbors use is motivated by certain property-related crimes: robbery rates and motor vehicle theft play an important role as well. But this analysis ignores several crucial parts of how Ring Neighbors operates. First, counties are a poor unit of spatial measurement; race and income can be highly mixed within entire counties, masking effects that we might be able to measure with a smaller unit of measurement. Second, to truly answer questions related to how Ring is used, the content of posts needs to be analyzed. We address both of these problems through a detailed case study of how Ring Neighbors has been used in Los Angeles, one of the most Ring-dense urban areas in the entire United States.

3 AN URBAN CASE STUDY: LOS ANGELES

Los Angeles was chosen as the case study because it is a large, racially and economically diverse city with relatively high Ring usage per-capita. As of 2019, Los Angeles is the second largest city in the United States, with nearly 4 million residents, of which 48.5% are Hispanic or Latino, 11.6% are Asian, and 8.9% are Black [5]. The city of Los Angeles was chosen instead of the Greater L.A. Area because of data availability

concerns: there are hundreds of municipalities and unincorporated areas in the L.A. metropolitan region, and data such as 311 calls or police crime reports are not available in all of these communities. In May 2019, the Los Angeles Police Department also became one of the earliest law enforcement agencies in the U.S. to partner with Ring, and the LAPD has used the platform extensively to collect video footage. Most notably, in the summer of 2020, the LAPD partnered with other local departments to create the "Safe L.A. Task Force" to surveil and prosecute Black Lives Matter protests, and repeatedly requested surveillance footage from Ring camera owners as part of this task force [8, 18]. Los Angeles, like many American cities, also has an extensive history of racial segregation and discriminatory housing policy, as well as a prominent culture of exclusive, suburban homeownership [15]. This history of racial gatekeeping enables us to study the ways Ring interacts with, and potentially exacerbates, existing patterns of exclusion within Los Angeles. Additionally, an author is from Los Angeles, which helps us better understand and contextualize findings.

One of our driving questions in examining Neighbors use in L.A. is whether Ring is used as a digital tool for racial gatekeeping and policing, particularly by white neighborhoods. Without analyzing demographic trends over time, which we leave to future work, there are two settings we identify as potentially precipitating white-led racial gatekeeping practices on the platform. First, white areas surrounded by other white neighborhoods—white "enclaves"—may feel a sense of racial anxiety and protection regarding their neighborhood, and a heightened paranoia of people perceived as outsiders. This might prompt higher posting rates, particularly if these posts primarily depict people of color. Second, we might expect to see white areas bordering more non-white neighborhoods attempt to surveil non-white people and use the platform, as Kurwa argues, to dictate the terms of potential racial integration [23]. We explore these questions both through a spatial modeling of ring posting rates and through our examination of post content below.

We also hypothesize that Ring usage in L.A. may have a relationship to other forms of incident reporting in the city, particularly 311 calls. Prior work has used 311 reporting rates as an estimate of civic engagement [26], which may correlate with the use of a platform like Neigbhors, particularly as a community safety tool. 311 calls have also shown to be driven by what Dan O'Brien refers to as "territoriality", in the sense of ownership and agency regarding the built environment a community occupies [34], and calls related to otherwise innocuous occurrences like "loud music" have been used as a proxy for inter-neighborhood conflict (i.e. residents call 311 as a way to harass and effectively police another community) [25, 26]. This makes 311 calling rates more interesting than just a potential control variable. Positive relationships with 311 calls would offer further support for the gatekeeping hypothesis.

Related to the gatekeeping hypothesis, examining L.A. also provides an opportunity to investigate the degree of *framing* that users on the Ring platform engage in. The core mechanism through which gatekeeping might occur on Ring is through framing subjects of videos as criminal or suspicious, regardless of the actual content shown and shared. Understanding if this pattern happens on the platform is crucial to characterizing its use and whether gatekeeping is a core function.

To answer these questions and further characterize Ring Neighbors in Los Angeles, we analyze how residents use the platform using three core methods. First, we use a spatial modeling approach similar to the one used at the national level to estimate the impact of different tract-level demographics and civic behavior on rates of Neighbors usage. The small size and relative demographic homogeneity of tracts in L.A. allows us to test more specific hypotheses related to racial gatekeeping and Neighbors use than in the national study. Second, we leverage an unsupervised machine learning technique called Structural Topic Modeling

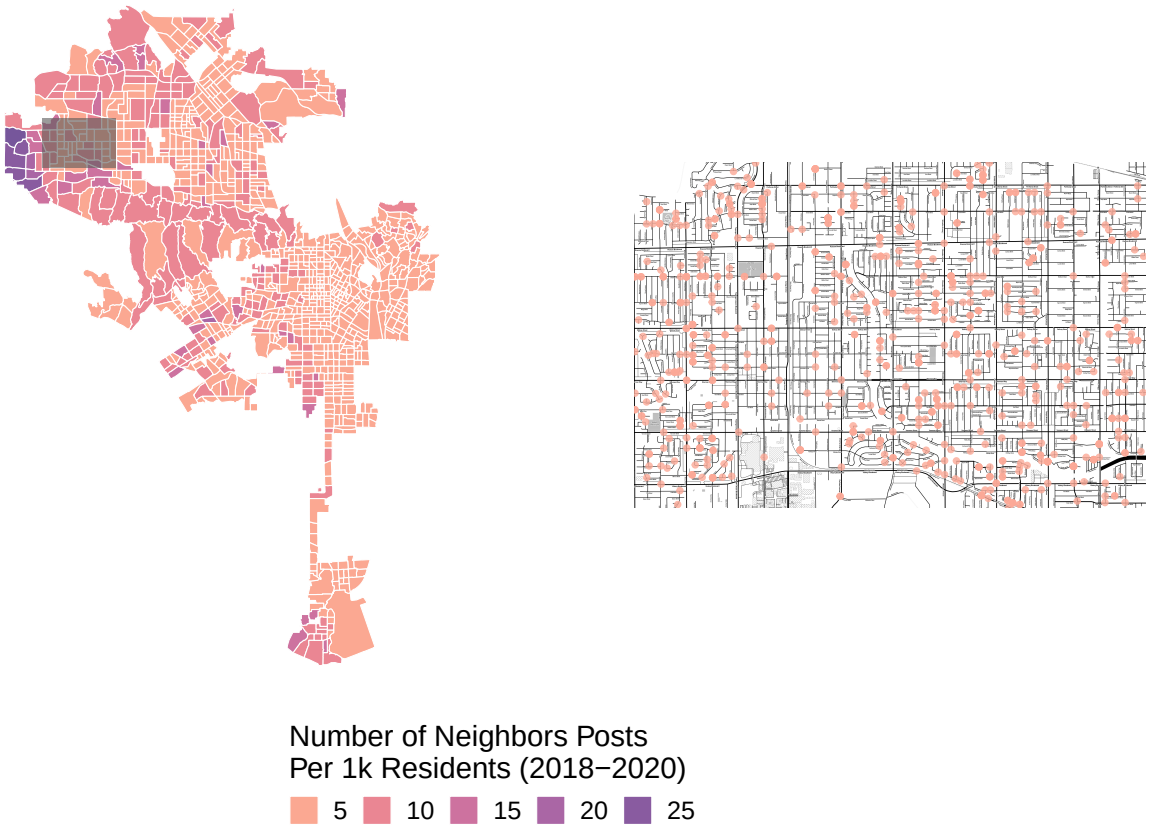

Fig. 4. Map of Neighbors alerts posted per 1,000 people in each census tract in Los Angeles from January 2018 - February 2020 (left), and dot map of street-level locations of Ring alerts from scraped data (right). Each dot represents a location recorded by Ring for a post on their platform, and are often the location of the camera that recorded uploaded content. Locations are pinned to the nearest street corner.

(STM) [41] to extract topics from the posts made in L.A, and further characterise these topics through a deep reading of each topics' most representative posts and a meta-categorization defined by the authors. Third, we augment our dataset with an experimental survey designed to both add structured coding data to posts and measure the degree to which post authors frame content on Ring as suspicious or criminal.

4 MODELING NEIGHBORS USE IN L.A.

4.1 Methods

As a first step towards understanding if Ring is used as a digital tool for racial gatekeeping and policing in Los Angeles, we construct a spatial model similar to the one used in Section 2 to explain Neighbors posting rates at the tract level across L.A. We use posts made in L.A. between Jan 1, 2018 and Feb 15, 2020, effectively measuring behavior on the platform from 2018 through the end of our observation period. We use posts from 2018 and later to match posting data to covariate data from the 2015-2018 ACS. We use crime reporting data from 2019 as reported by the city of Los Angeles, and merge reporting rates into five

main categories. We split theft into two categories: theft, which includes personal theft, and property/house theft, which here includes burglary and robbery of a home.

We include standard demographic covariates, such as median household income and basic racial demographics, and we also construct variables to measure constructs related to racial gatekeeping by white neighborhoods. Racial demographics represent the percentage of a tract's population that identifies as that race *alone*, and we choose to measure white population share as the share of *non-hispanic white* residents in a tract to account for possible racial dynamics between hispanic communities in L.A. and majority-white communities.

To more explicitly operationalize racial gatekeeping in our model, we create a new variable for each tract, `pct nonwhite neighbors`, which represents the percentage of neighboring tracts that are not majority non-Hispanic white as measured by the 2015-2018 US ACS. We define "neighbors" as any tracts that share a part of any border. We also construct a dummy variable, `maj. white`, indicating if a tract is majority non-Hispanic white. We add the interaction of these two terms into our spatial model of L.A. usage, allowing us to estimate the impact of a tract being a white "enclave" as well as the impact of higher rates of non-white neighbors in majority-white tracts. To control for civic engagement and to test our hypotheses related to 311 calls, we include tract-level rates of 311 reporting from 2019 as covariates in our model.

4.2 Findings

Figure 5 shows spatial modeling results for tract-level Ring Neighbors usage in Los Angeles. Here, we report direct impacts—the impact of a covariate change on that tract itself—because some of our covariates are spatial themselves.

4.2.1 White 'Enclaves' Post More Than Other Tracts. The impact of a tract being majority-white with no non-white neighbors—our "white enclave" term—stands out in our model results. If a tract is one of L.A.'s 136 white enclave tracts, its residents post over 18 more posts per 1,000 residents than others, controlling for our other covariates. It is worth remarking on the insignificant impact of the non-Hispanic white population share. It is not true that all majority white communities are more likely to use Neighbors. Instead, majority white areas, surrounded by other white neighborhoods, are associated with this huge jump in posting.

4.2.2 311 Calling Rates are Positively Associated With Neighbors Use. We also see that 311 calling rates have some significant impacts in our model, particularly with majority-white tracts. Calling rates for the categories 'cleaning', 'homeless', and 'feedback' are all positively associated with Neighbors posting rates, but only in white tracts. In particular, correlations with 311 homeless calls in white majority tracts lends validity to our racial gatekeeping hypothesis; in Los Angeles, the unhoused population is disproportionately Black [1] and 311 calls reporting homeless people or encampments are often followed by sweeps and police presence [38]. As such, positive correlations between Ring posting rates and 311 homeless calls in white majority neighborhoods suggests a shared tendency to police and report "unwanted" neighborhood members.

4.2.3 Race and Homeownership. These spatial modeling results indicate that race and homeownership plays a clear role in Ring posting rates: white enclaves, or white majority census tracts bordering only other white tracts, post on Ring at significantly higher rates and rates of property ownership are positively correlated with Ring usage. Additionally, correlations between 311 calling rates, especially to sweep homeless

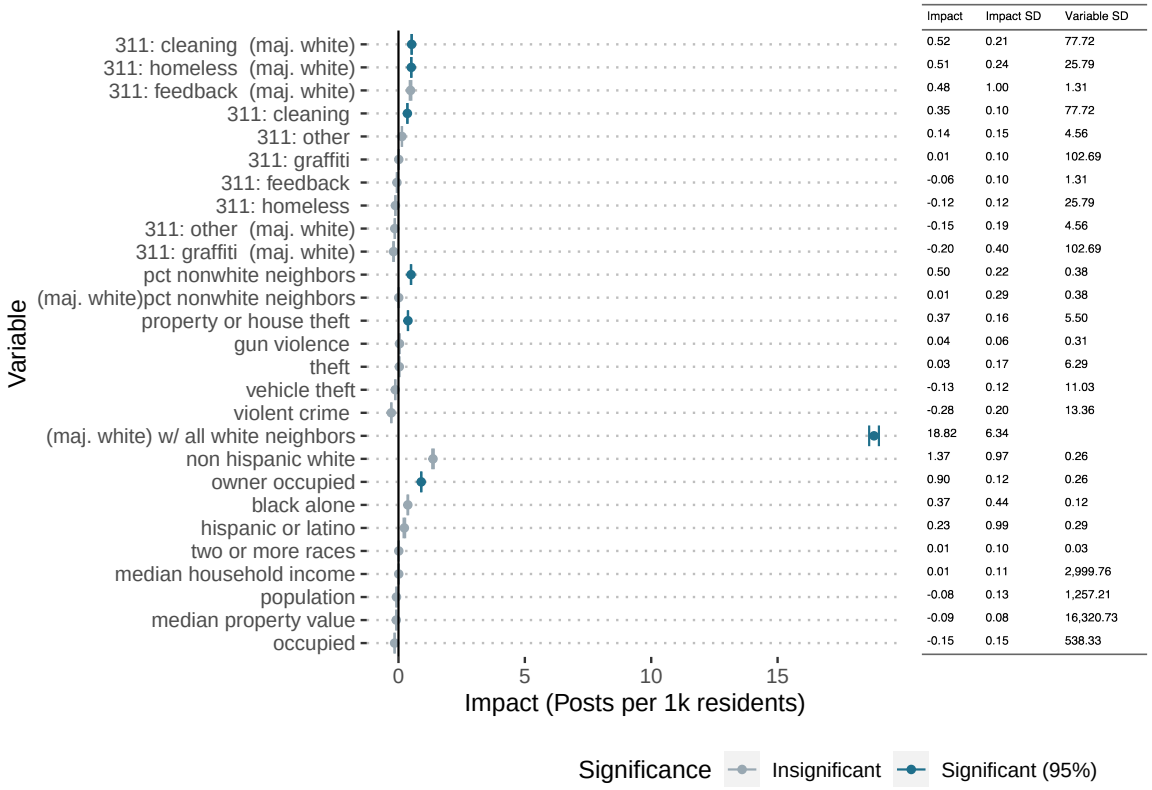

Fig. 5. Spatial Autoregressive Lag model results for L.A. tracts. All variables are centered and standardized before modeling. Nagelkerke pseudo-R-squared is 0.678. The Impact column refers to the average direct and indirect impacts of one S.D. of change in that variable on the per-capita number of Ring Neighbors posts in a tract in 2019. Bars shown are estimated standard errors. Median household income and property values are in US dollars. 311 reporting variables are in number of calls from a tract. The variable "pct nonwhite neighbors" refers to the percentage of neighboring tracts whose residents are majority non-white. Crime variables, like violent crime and vehicle theft, are in incidences per 1000 people in 2019. Variables labeled with "(maj. white)" refer to interaction terms with a logical variable indicating if a tract is majority non-Hispanic white. Variable standard deviation is omitted for logical variables.

encampments, and Ring posting rates in white enclaves provides preliminary evidence towards gatekeeping tendencies.

5 TOPIC MODELING AND ANALYSIS OF LA NEIGHBORS POSTS

While our spatial analysis provides some evidence consistent with theories of Ring as a racial gatekeeping tool, it does not answer how users might perform this gatekeeping or how the platform is used more generally in LA. As a case study in how Ring Neighbors is used in a major US city, and to test our hypotheses about its use as a gatekeeping tool, we first use structural topic modeling (STM) to extract general topics from posts on the platform. Then, through a deep reading of the most "representative" posts for each topic, we code topics into "meta-categories" and qualitatively validate the STM output. The STM regression framework then allows us to estimate the association between tract covariates such as race or income with

the frequencies that tracts post about certain topics. This approach allows us to discover some overall patterns between neighborhood demographics and post topics.

These methods only examine post text and location, however. Video content is core to how Ring is used, and so another approach is needed to answer basic questions about post content, particularly as related to topics. For example, the second-most frequent topic we discover (Topic 2, see Figure 6) appears to mostly include posts that depict strangers knocking on doors. Posts the authors reviewed in this category largely claim more suspicion than we judged to be warranted by the posted videos. Many posts in this topic (and others) described what we deemed to be innocuous activity as heavily suspicious or "shady", a framing we suspect to be pervasive on the platform.

However, our positionality as authors makes an unbiased review of such patterns difficult. We arrive at analyzing posts having already framed Ring as being within a racialized carceral logic. Rather than only conduct a deep qualitative reading from this positionality, we opt to generate a more unbiased, or at least representative, view of content on the platform through a quantitative analysis backed by our own readings of posts.

To do this, we developed a carefully crafted survey designed to both code post content and isolate perceptions of a post's video from the subjective framing added by Ring users in post titles and descriptions. We deployed our survey on Amazon Mechanical Turk, and randomly assigned workers to one of two conditions. Workers in each condition were asked to answer questions, such as rating the severity of any shown crime, for a series of randomly selected posts. In the first condition, workers answered questions about full posts, including its title and post text. In the other, workers only saw the post's associated video. By comparing survey answers in each condition, this approach allows us to compare, for example, the overall rates of crime claimed in a topic to the rates of crime shown in the videos in that topic.

Together, the two approaches described above allow us to report broad statistical patterns of how Neighbors is used in LA, some characteristics of the posts themselves, and a measurement of the level of framing users employ on the platform to cast video content as suspicious or criminal. We focus on answering the following questions:

- **RQ1. To what extent do Ring users frame innocuous content on the platform as suspicious?**
- **RQ2. How often do Ring users claim criminal activity, and how often is this activity actually portrayed in posted videos?**
- **RQ3. Are there differences between the posting behaviors of our three main neighborhood categories: "white enclaves", "nonwhite" and "white with nonwhite neighbors"?**
- **RQ4. Are there other major differences between posting behaviors explained by other demographics we found important in our regressions, such as owner occupancy, 311 calls, or theft rates?**
- **RQ5. What are the racial characteristics of subjects filmed by posters?**

This section is organized in three main parts. First, we report the methods used for the grounded computational theory analysis, including the topic analysis, the authors' "meta category" topic coding, and the regression method. Then, we detail our experimental survey design, detailing the survey and how we isolate characteristics of post videos from posts as a whole. Finally, drawing on this collection of methods, we discuss our results, finding that the way users frame posts has a significant impact on the perceived

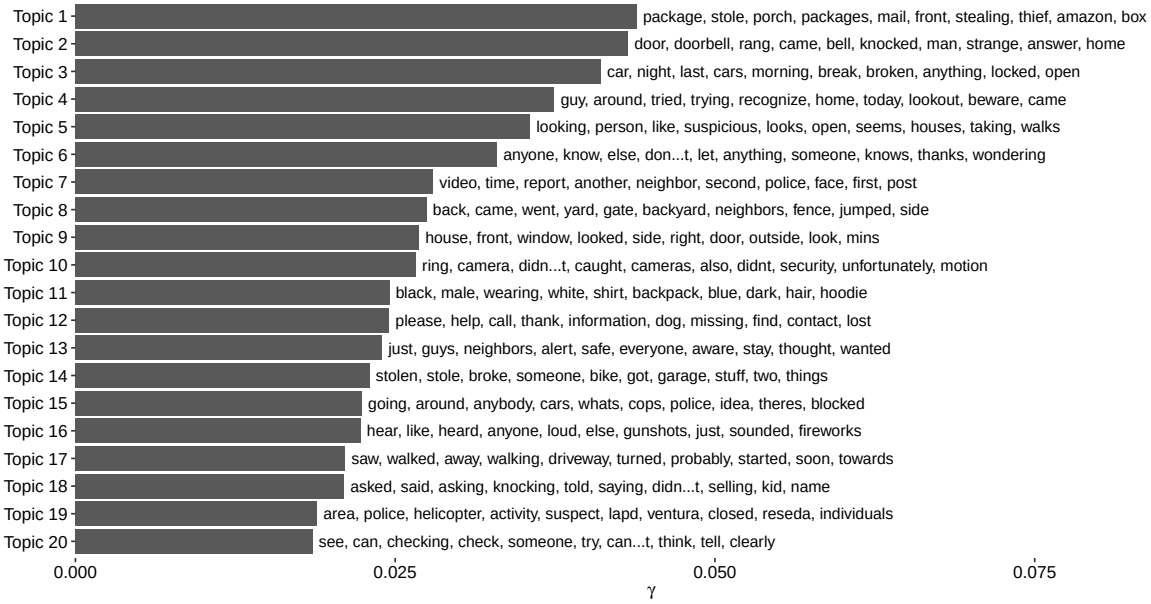

Fig. 6. Top 20 topics and their top 10 terms extracted from all public Ring Neighbors posts in Los Angeles from Jan. 2018 through Feb. 2020. The x-axis represents the overall prevalence of that topic.

suspiciousness of videos, that users claim criminal activity more often than it actually appears in posts, and demonstrating that race explains significant differences in posting rates between neighborhoods

5.1 Methods

5.1.1 Topic Modeling and STM Regression. Each post on the Neighbors platform contains a title and post text, and an associated image or video if it is a media post. Our corpus for topic modeling consists of the post text from all posts made during our observation period on the Neighbors platform in Los Angeles. We use only post text for convenience of analysis and data consistency: many post title terms appear in post text, and post text is usually written in whole sentences. We choose not to stem our documents, following the recommendations of Schofield et. al. [44], and our processed corpus after removing stop-words, numbers, and punctuation, consists of 8,916 documents with 5,613 unique terms. We train nine different STM models, with K ranging from $K = 10$ to $K = 90$ using the stm package for R [41] with spectral initialization and 50 EM iterations. To determine the number of topics for our final STM model for further analysis, we use a heuristic combination of diagnostics. We choose a K that provides a balance between semantic coherence, model residuals, and exclusivity [30], resulting in a model with $K = 60$. Topics are labeled in order of frequency (i.e. Topic 1 is the most frequent topic).

The STM provides two central metrics. For each topic, we report its γ, which represents its general prevalence in our corpus. For example Topic 1 is our most prevalent topic, with $\gamma = 0.043$, indicating that Topic 1 represents 4.3% of all content we analyzed. Each post also then has an associated β value for each topic, indicating the percentage of its content that the STM attributes to a given topic.

The most common terms and their associated prevalence (how often they appear in that topic) for the top 20 topics inferred by STM in our corpus is shown in Figure 6. The most prevalent topic in our corpus, Topic 1, seems to be specifically about package theft. Topic 2 appears to be about strangers knocking on doors or ringing doorbells. These are sensible and serve as a helpful sanity-check for our modeling. Our topic analysis also reveals different ways users use the Neighbors platform. Topic 6 shows what appears to be community-oriented posts, with language indicating posters asking questions or warning other Neighbors users about some event. Topic 11 may reveal some racialized use of the platform: its top terms are "black" and "male".

The STM framework also allows us to regress the prevalence of topic usage on the same tract-level covariates we used in Section 4, providing some insight into tract-level differences in topic use. STM regression results are shown in Table 2. Estimates are shown as the average treatment effect of a unit change in that covariate on the γ value (overall prevalence) of a topic. Units here are presented not as S.D's, as in Section 4, but instead as natural units. For example, a 1% increase in the share of owner-occupied units corresponds to a 0.4% increase in posting rates in Topic 2 (strangers knocking on doors).

Table 1. Topic meta-categories coded by the authors.

Category	Topics
Community	6, 12, 15, 16, 19, 21, 27, 28
Crime	1, 3, 8, 14, 29, 30
Strangers	2, 4, 5, 7, 9, 11, 17, 18, 22, 26

5.1.2 Author Coding of Topics and Posts. To further categorize posts on the platform, and validate the topic modeling results, we extract the 10 posts with the highest β for each topic, and manually examine each post. This sampling method was chosen because it provides an unbiased sampling of the "most representative" posts for each topic. To categorize posts, two of the authors independently annotated each post and video with qualitative notes (see Appendix B), and then extracted meta-themes across topics, such as "Theft and Crime" or "Suspicious Strangers". Authors then compared annotations and came to consensus on two kinds of labels: first, a brief, descriptive label for each topic (shown in Appendix B) and a broad meta-category label for each topic. Authors decided on three meta-categories that describe most topics in our corpus: "Suspicious Strangers", "Crime", and "Community Safety". Table 1 shows the topic membership for each of these meta-categories. For some validation of the authors' meta-categorizations, see Section 5.2.1.

5.1.3 Crowdsourced Coding and Experimental Survey. To further characterize each topic, we designed a survey to both code post content and to quantify the amount of "spin" or framing users might use in casting a video as suspicious or criminal. For each topic, we extract the 20 posts that are most representative of that topic (highest β) that have an associated video, resulting in 600 "representative" posts in total. We designed our survey in Qualtrics and paid 613 workers to participate in our survey through Amazon Mechanical Turk. Following recommendations from recent work examining workers on Mechanical Turk, we select US-only workers who have done less than fifty total studies [42].

For each post, workers were asked to answer between five and seven questions (a full survey design, including full question text, can be found in the Appendix):

1. **Race**: Race of each person portrayed in a video (count of each racial group). We collect data on the number of people coders perceived as White, Asian, Black, Hispanic, and Other.
2. **Claimed Crime**: Whether the author claimed any criminal activity occurred in the post (binary yes/no).
    1. If answered "yes": **Severity (Claimed Crime)**: The severity of any claimed criminal activity (Likert, "Very Minor" to "Very Severe").
3. **Police**: Whether the post mentions contacting the police (binary yes/no).
4. **Shown Crime**: Whether the post video portrays criminal activity (binary yes/no).
    1. If answered "yes": **Severity (Shown Crime)**: The severity of any portrayed criminal activity in the video alone (Likert, "Very Minor" to "Very Severe").
5. **Suspicion**: The suspiciousness of activity portrayed in the video (Likert, "Very Innocuous" to "Very Suspicious").

Each worker was assigned seven randomly selected posts, and was assigned randomly to a control or treatment condition. In the control ("Video and Text") condition, participants saw the video, its title, and its post text. In the treatment ("Video Only") condition, participants saw only the video associated with the post. The survey questions were the same across both conditions. We exclude any ratings from workers who fail basic attention checks, as detailed in the Appendix, resulting in a total of 357 respondents and 13 posts that received no valid responses. Because post assignment to respondents was randomized, not all posts had the same number of responses. On average, posts were examined by 4.2 respondents in each treatment group, with a standard deviation of 1.6.

Below, we use the results from this experimental survey in two key ways. First, we use responses to various survey questions as post codings to answer basic questions about post content. For these basic questions, we use responses to questions 2-4 from respondents in the "Video and Text" condition, as these respondents' answers reflect their judgment of the entire post, including text.

When discussing results, we refer below to the percentage of posts where crime was "claimed" and when crime was "shown". To compute the number of posts where a crime was claimed, we only use responses from workers in the control ("Video and Text") condition, because they were the only group that saw post text. Similarly, when reporting "shown" crime, we use only responses from the treatment ("Video Only") condition. In each case, we only include posts (1) with at least two raters from the relevant treatment group, and (2) majority agreement.

We also refer to the number of people portrayed in a post's video when discussing posts and topics. To estimate the number of people of different (percieved) races portrayed in posts, we use responses to question 1 from respondents in the "Video Only" condition to reduce racial bias from post text.

Table 2. Significant (>95% level) STM regression results. Empty cells indicate no significant effect. Crime reporting rates are shown in incidences per 1,000 people. Property value is in tens of thousands of dollars. Percent nonwhite neighbors and owner occupation rates are in 1% units. For example, the result for owner occupied rates on Topic 2 can be read as: 'A 1% increase in owner occupancy rates is associated with a .43% increase in Topic 2 posting rates.' Results where the associated change in topic prevalence is less than 0.05% are omitted for convenience, but a full regression table can be found in Appendix C.

term	owner occupied	(maj. white)pct nonwhite neighbors	(maj. white) w/ all white neighbors	pct nonwhite neighbors
Topic 2	0.43% (0.13%)***	0.4% (0.2%)*		
Topic 5			-30% (12%)*	
Topic 7				-0.44% (0.16%)**
Topic 8	-0.39% (0.12%)**			
Topic 9	0.4% (0.087%)***			
Topic 13	0.39% (0.11%)***			
Topic 14	-0.74% (0.18%)***			
Topic 15			54% (24%)*	
Topic 17	0.45% (0.14%)**	0.53% (0.22%)*	-43% (15%)**	
Topic 18	0.9% (0.25%)***			

5.2 Findings

5.2.1 Post Meta-Categories. To organize and make sense of the topics generated by STM, we perform a qualitative deep reading of representative posts within each topic with the goal of organizing topics into "meta categories": groups of topics that are thematically similar. Through this qualitative reading we find that content on Ring can be grouped into three main meta-categories. First, the most common kind of post describes, portrays, or discusses suspicious strangers or activity. The second most common kind of post can be described as reports of crime. These posts often portray straightforward accounts of package theft, trespassing, or car break-ins and theft. Third, and least commonly, we find that Neighbors is also used as a community messaging board, where users ask other people about police action, lost pets, or other safety topics. This results in three meta-categories that we assign each topic to: "Suspicious Strangers", "Crime", and "Community Safety". A description and some qualitative readings of posts from each of these meta-categories is below.

One objection to this categorization may be that posts or topics about suspicious strangers may overlap significantly with posts or topics that mention crime or criminality. To validate our meta-categorization, we leverage post codings from our survey. Figure 7 shows the percentage of posts where a crime is shown or claimed between our "strangers" and "crime" meta-categories. Overall, posts the authors coded as belonging to the "crime" meta-category have higher rates of shown crime. On average, we also find that

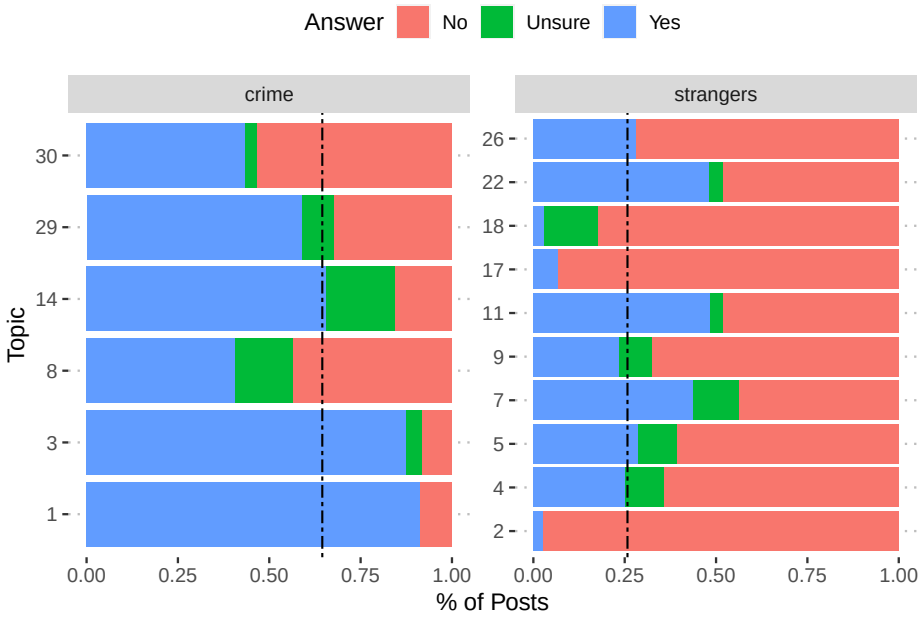

Fig. 7. Percentage of posts in each meta-category where a crime is shown or where a poster claims a crime occurred across topics between our 'suspicious strangers' and 'claimed crime' meta-categories. Dotted line in each plot represents the mean percent of posts across topics that are labeled as either claiming a crime or portraying one.

only 40% of posts in our community safety meta-category depict people, in comparison to 75% in our 'crime' meta-category and 87% in our 'strangers' meta-category, validating that posts in these topics are less likely to be about individuals, supporting the authors' categorization.

5.2.2 Category 1: Strangers. The most prevalent meta-category of posts on Ring in LA involves videos of people, often referred to as strangers or as suspicious, doing a broad range of activities. Posters also use language which assigns suspicious or criminal intent to innocuous or non-criminal activity, a theme we return to in Section 5.2.5.

For instance, Topic 2 is mostly about people—often described as "strangers"—knocking on doors. Topic 2 appears nearly as frequent as the top topic in our corpus, which is about stolen packages (γ=4.4% vs γ=4.3%). In many cases, the post authors describe the act of knocking on a door as suspicious. For instance, the most representative post in Topic 2 (β=62.6%) with a video is titled "Guy knocking" and shows a person walk up to the poster's door and knock hard for a few seconds. That a user chose to upload this activity without further context is notable. Another poster describes a person knocking on the door as a "shady guy" and writes "who is this guy and why does he knock on the door instead of ringing the doorbell?", suggesting that it is inherently noteworthy or suspicious to knock on a door. 80% of the top 20 posts in Topic 2 were rated by coders as portraying "somewhat innocuous", "innocuous", or "very innocuous" activity, supporting the idea that most activity filmed in Topic 2 is generally considered innocuous.

Additionally, posts in Topic 5 (γ=3.6%) are reports of "suspicious people" who, according to the post authors, look like they want to commit a crime. For instance, the most representative post in Topic 5

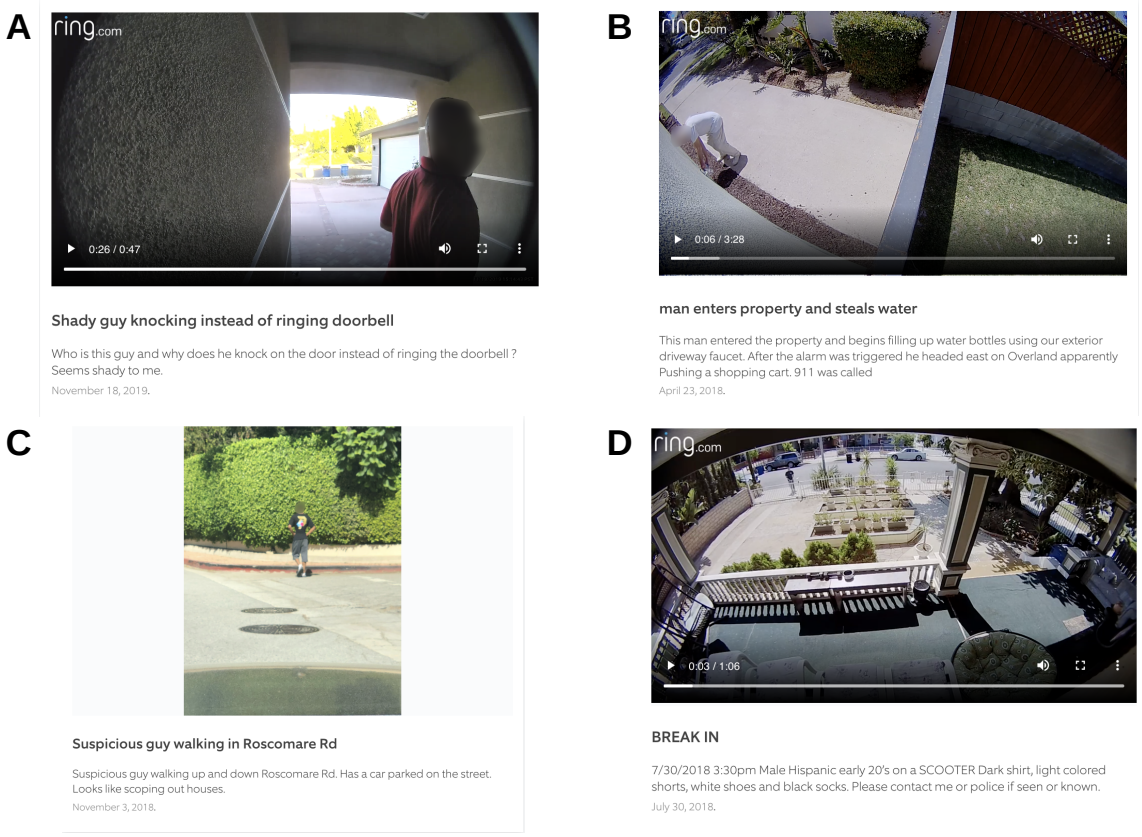

Fig. 8. A: Neighbors posts often show innocuous activity (knocking on a door) but assert it is suspicious. B: An example of how Ring users criminalize and involve the police over minor incidents. C: Many Neighbors posts show ambiguous activity but label it—without further context—as suspicious or, often, criminal. D: An example of how racial language is used on Ring. Here, it is hard to determine the filmed person's race but the post author describes them as "Hispanic". Screenshots taken by the authors.

(β=46.0%) shows a young adult walking to the entrance of the post author's condo, looking around for a few seconds, and walking away. In the description, the author writes "This suspicious character looks like he's up to no good. Possibly scoping out condos while owners are away enjoying the holidays. Keep an eye out for this guy!" (see Appendix A #10). Other posts in Topic 5 describe a "Suspicious guy walking up and down [a road]... Looks like scoping out houses" or a "woman on my porch at 2:51am... [who] looks like she was looking for something to steal" (see Figure 8C and Appendix A #11, 12).

Another example of how Ring is used to post about activity users consider "suspicious" is Topic 17 (γ=2.1%), which mostly reports strangers such as potential package thieves or people supposedly pretending to inspect homes. Posts in Topic 17 use terms related to walking such as "saw, walked, away, walking, driveway", and also discuss activity that the poster believes is suspicious. For instance, the most representative post in Topic 17 with a video (β=31.1%) is titled "Same Creepy Homeless Guy!" and depicts a person at night walking up to the poster's door and pacing around and adjusting their clothing on the porch for thirty

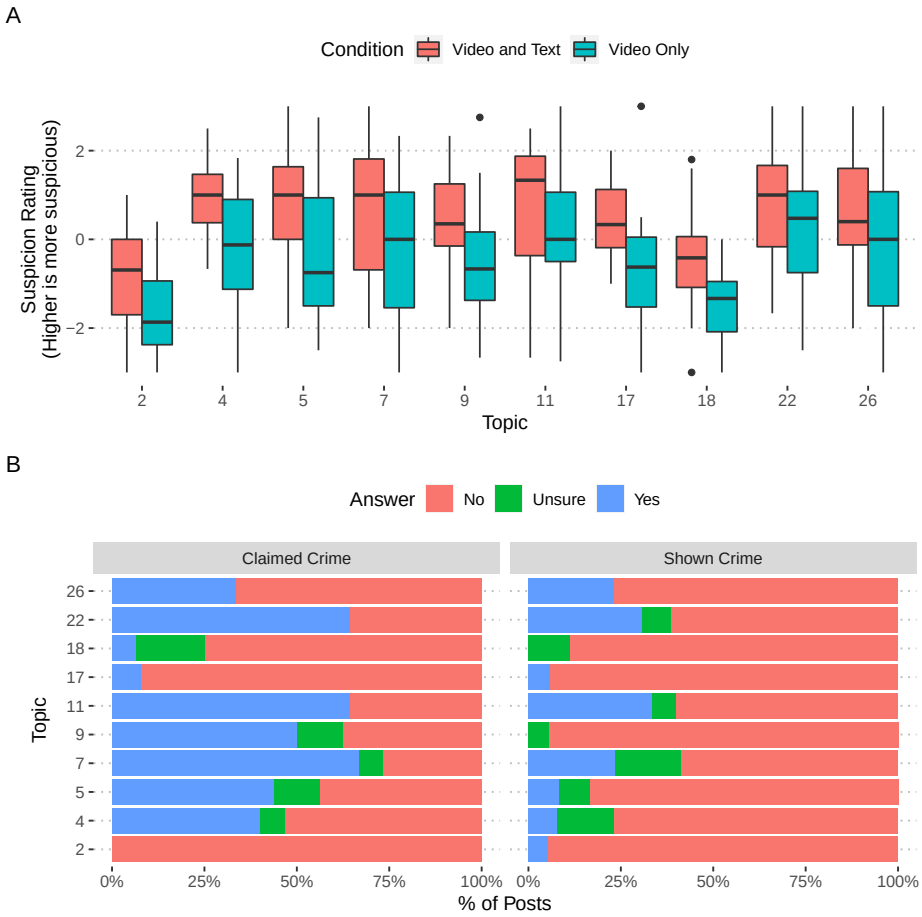

Fig. 9. Experiment and coding results for 'stranger' meta-category. Figure A shows the difference in suspicion ratings between respondents in our 'video only' and 'video and text' conditions for 'stranger' topics. Figure B shows the percentage of posts in each topic that were labeled as claiming a crime (left) and showing a crime on video (right). Posts are only included in Figure B if a majority of annotators agreed on a label.

seconds before walking away. When shown only the video for this post, all raters labeled it as "innocuous" or "somewhat innocuous". However, when shown both the video and text, all raters labeled it as "suspicious". Thus, this post is an example of an instance where posters record an innocuous activity, and cast it as suspicious.

Other examples of users on Ring posting about "suspicious" activity are Topics 4, 7, 9, 18, and 26. Topic 4 (γ=3.7%) mostly reports strangers looking around the author's property and sometimes trespassing; Topic 7 (γ=2.8%) contains reports of suspicious or criminal activity, and many posts reference filing police reports or sharing footage; Topic 9 (γ=2.7%) mostly reports "suspicious vehicles", such as unknown cars parked outside where post authors often accuse these drivers of inspecting houses for a potential burglary; Topic 18 (γ=2.1%) mostly posts about suspicious solicitors, often young adults who are asking for money in what post

authors label scams; and Topic 26 (γ=1.5%) describes a variety of strangers, such as a "possible package thief", "suspicious gentleman", or a suspicious "white cadillac".

In some cases, Ring posters use racialized language when describing "suspicious" people or criminal activity. Posts in Topic 11 (γ=2.5%) often claim criminal activity is occuring (64% of posts), and posters will also report suspicious people like a "prowler" or "5 suspicious males casing homes". The most representative terms in Topic 11 are "black, male, wearing, white, shirt, backpack, blue, dark, hair, hoodie", which suggest that posters in this topic often mention race. Though terms like "white" and "dark" can refer to color descriptions, a qualitative reading of the posts reveals that these terms are used in racialized ways. For instance, the most representative post with a video in Topic 11 (β=75.4%) is titled "Another construction thief" and shows a person wearing a traffic vest walking around the Ring owner's parking lot. Though raters determined that the video showed "somewhat suspicious" activity, they were "unsure" that a crime was depicted, which is consistent with the pattern in this meta-category of assigning criminal intent to non-criminal activity. Notably, the post author includes a detailed description of the filmed person, using racialized language ("African American male", "pants, sagged like a thug"). Other posts also include racial descriptions ("Male Hispanic", "Male black 5-8", "6 foot white male"), or reference physical attributes ("brown hair", "bald", "wearing a black cap").

In other cases, Ring users will also call the police in response to activity they believe to be suspicious. For instance, Topic 22 (γ=1.8%) contains reports of petty crime and neighborhood commotion such as a fight between "3 Young Men ... w/ Bricks" (see Appendix A #5) as well as posts about "suspicious" activity like unruly teenagers, strangers washing their hair in the front water spout at night, and a person "walking down [a] street which is a dead end" who made some verbal threats (see Appendix A #7, 8, 9). A majority of posts in Topic 22 coded by human raters (14/20) contained references to calling the police, but of the posts where there was a reference to calling the police, only 4/14 posts (28.6%) were coded by raters as portraying clear criminal activity.

5.2.3 Category 2: Crime. With regards to RQ2, we find that the second most frequent type of Ring post in Los Angeles is reports of crimes such as package theft, break-ins, and burglaries. These posts are often straightforward, informal descriptions of a crime which affected the post author and usually include a video or a photo.

The main type of crime reported on Ring is theft: the most discussed topic on Ring, Topic 1 (γ=4.4%), contains mostly reports and videos of package theft; Topic 3 (γ=4.1%) primarily reports car break-ins; Topic 14 (γ=2.3%) is mostly about stolen bikes and theft from automobiles; and Topic 30 (γ=1.5%) mostly reports theft and vandalism, often written in Spanish. As shown in Figure 10, a majority of posts in these topics claim that criminal activity is occurring and raters in Topics 1 and 3 (but not 14 or 30) also agree that crime is being shown in the majority of videos. Moreover, raters largely agree with posters on the severity of crime occuring. This supports the interpretation that posts in Topics 1, 3, and to a lesser extent 14 and 30 are straightforward reports of actual criminal activity.

Another common type of crime reported on Ring is trespassing: Topic 8 (γ=2.7%) mostly describes strangers breaking into backyards, often by climbing over fences or opening gates. 43.75% of posts in Topic 8 claim criminal activity is occurring, with raters determining that 37.5% of posts actually show criminal activity (see Figure 10C). This indicates that Topic 8 also contains many posts which are straightforward

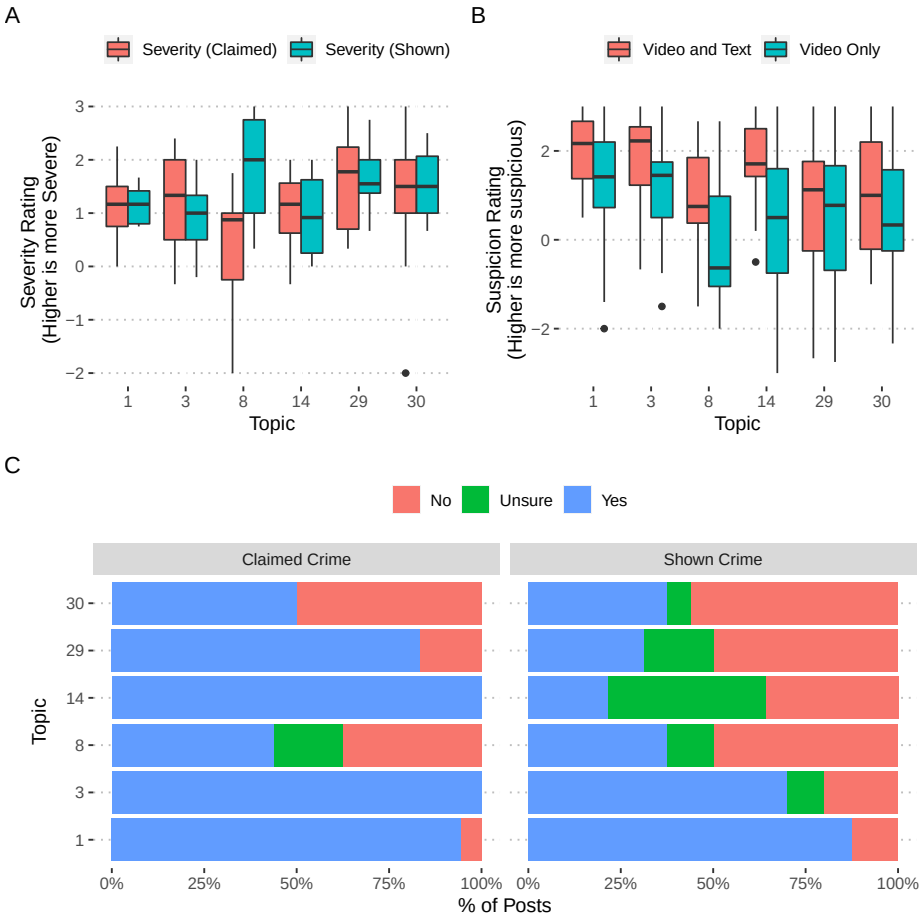

Fig. 10. Experiment and coding results for 'crime' meta-category. Figure A shows the difference in severity between crimes that were claimed by posters and crimes shown in videos. Figure B shows suspicion ratings across topics between full posts and videos only. Figure C, below, shows the percentage of posts where raters in the 'video and text' condition agreed that the post author claimed a crime was committed (left), and the percentage of posts where raters in the 'video only' condition agreed there was an actual crime committed on video (right). For example, while all posts in Topic 14 are rated as claiming criminal activity, under 25 percent of posts are labeled as actually depicting a criminal act.

reports of criminal activity. Interestingly, Topic 8 is the only topic where raters in the "video only" condition generally rated the shown crime more severe than those who also read the poster's framing.

Ring users also post about burglary, though it is less common: Topic 29 (γ=1.5%) mostly reports burglaries. As shown in Figure 10C, a majority (83.3%) of posts in Topic 29 claim criminal activity is occuring, though a much smaller percentage of raters (31.25%) agreed that criminal activity is actually shown in the videos. However, raters largely agree with posters on the severity of crime occuring.

5.2.4 Category 3: Community Safety. The last main usage of Ring in Los Angeles is as a community messaging board to ask questions about neighborhood commotion, ask for help finding lost pets, or discuss other safety

topics. This usage is consistent with how Ring frames itself as a local social network where neighbors can work together and share information.

Many posts on Ring are questions asking other Neighbors users to help identify commotion or unknown activity in the neighborhood. For instance, Topic 6 (γ=3.3%) contains questions about neighborhood activity such as nearby police and lost pets; Topic 15 (γ=2.2%) mostly contains questions about commotion such as road blockages, smoke, or police in the area; posts in Topic 16 (γ=2.2%) often ask about loud noises such as gunshots, fireworks, and explosions; and Topic 27 (γ=1.5%) mostly contains questions asking about nearby helicopter activity.

In addition to asking questions, users on Ring also make safety related announcements. For instance, Topic 19 (γ=1.9%) mostly reports local police activity and includes announcements written by law enforcement; Topic 21 (γ=1.8%), contains mostly announcements about a variety of safety related topics like gunshots or police nearby; and Topic 28 (γ=1.5%) primarily contains posts instructing other users to be careful or aware of dangerous wildlife, aggressive pets, or other animal related topics.

Additionally, another common use case of Ring is to ask for help finding lost pets or missing relatives: Topic 12 (γ=2.5%) mostly contains posts notifying neighbors about lost pets or, less commonly, missing family relatives.

5.2.5 Posters Frame Some Content as Suspicious or Criminal, Often Without Evidence. As discussed in Section 5.2.1, the primary type of content on Ring is reports of activity a poster believes to be suspicious: anything from knocking on doors, "scoping out condos", or walking in an unusual way. Using human coders and a randomized controlled trial where respondents are shown Ring posts with ("Video and Text") or without post text ("Video Only"), we also find evidence that posters on Ring frame filmed activity using language which assigns suspicion or criminal intent to innocuous or non-criminal activity for topics about "suspicious strangers", which helps answer RQ1 and RQ2.

For the "suspicious strangers" meta-category, raters in the "Video Only" condition rate these posts as more innocuous than raters in the "Video and Text" condition (see Figure 9 A). This implies that Ring posters use language which causes respondents to view the same filmed activity as more suspicious. Additionally, for topics about "suspicious strangers", Ring posters often claim criminal activity is occurring even when human coders do not believe that the filmed activity depicts a crime (see Figure 9 B).

For instance, in Topic 5, raters in the "Video Only" condition label posts as innocuous; however, in the "Video and Text" condition, the valence switches and raters label the same posts as highly suspicious. Moreover, while around 40% of posters in Topic 5 claim criminal activity is occuring, under 25% of posts are labeled by raters as actually depicting a criminal act. This is also a pattern across Topic 17: when assigned to the "Video Only" condition, raters on average label posts in Topic 17 as more innocuous, but when assigned to the "Video and Text" condition, raters label these posts as more suspicious. Additionally, in 8/20 of the most representative posts in Topic 9, authors claimed that there was criminal activity; however, in only two of these eight posts did raters agree that there was clear criminal activity.

5.2.6 Differences Appear When Crime is Claimed, but not Shown on Video. Are the crimes claimed by posters rated as more severe than the crimes depicted on video alone? To answer this question, we split posts into two categories. The first category (1) consists of posts where the author claimed a crime, but where raters

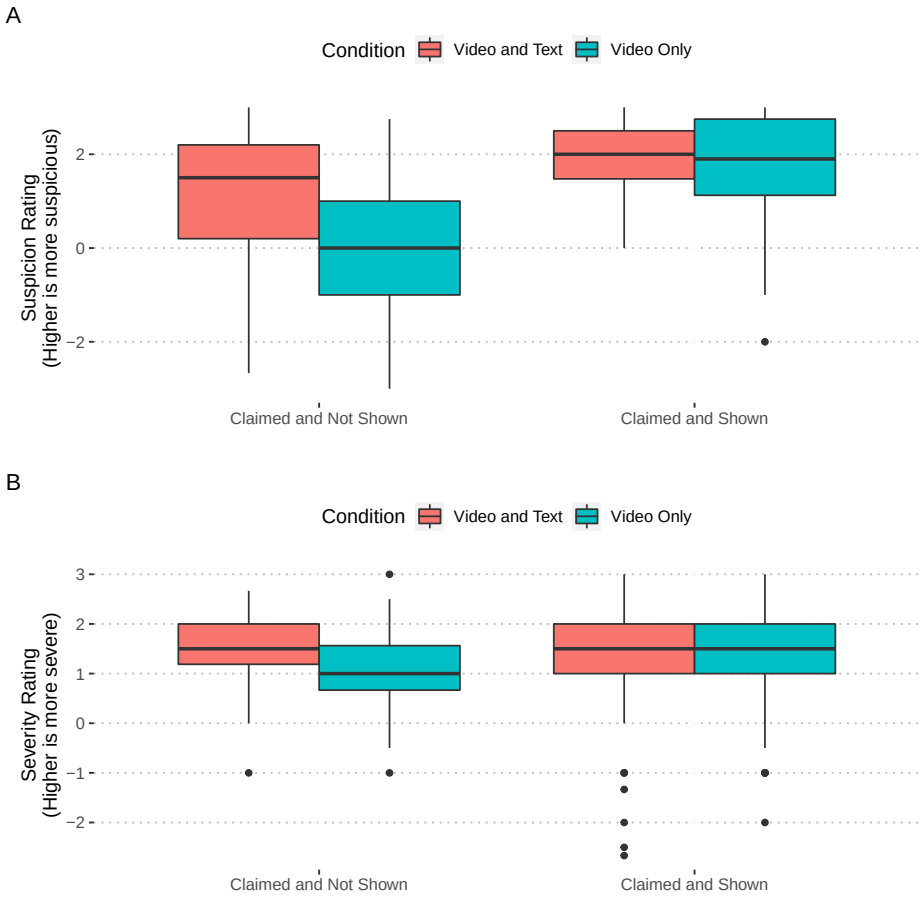

Fig. 11. Suspicion (A) and Crime Severity (B) ratings for posts where a crime was claimed. The left column on each plot represents posts where a crime was claimed but not shown (i.e., raters coded the post as not showing a crime). The right column shows posts where a crime was both claimed and coded as being shown on video. Posts where a crime is claimed but not shown are the source of differential ratings of both suspiciousness and severity.

agreed no crime was depicted in the post video. The second category (2) consists of posts where a crime was both claimed by the poster and shown in the associated video.

For each of these groups, Figure 10B shows the difference in rated severity between raters that saw the full post vs only the post video. When a crime is claimed but not shown (1), the claimed crime is more severe than the shown one, but when a crime is both claimed *and* shown (2), raters in both conditions agree. The same phenomenon is true for suspicion ratings, shown in Figure 10A. This shows that posts that claim crime without showing clear evidence actively frame activity on video as criminal and suspicious, impacting raters' perceptions.

5.2.7 White Neighborhoods with More Nonwhite Neighbors Post More in Some Topics About Strangers. Majority-white tracts with more nonwhite neighbors tend to post more in some topics about strangers than non-white tracts or "white enclave" tracts. For majority white tracts, a 10% increase in the share of

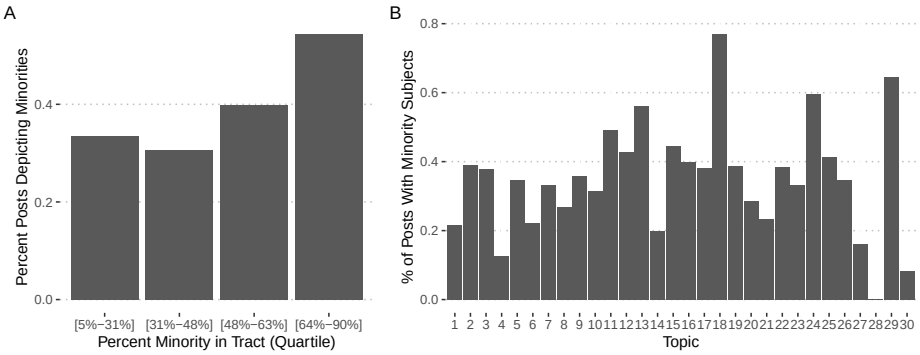

Fig. 12. A: Percent of posts depicting minorities by percent of minority residents in a tract. Each bar represents one quarter of all tracts. B: Percent of posts that depict minority subjects in each topic.

neighboring non-white tracts corresponds to a 4% increase in Topic 2's prevalence, and a 5.3% increase in Topic 17, both stranger-centric topics with high levels of "spin". Meanwhile, white 'enclaves' post 43% less about Topic 17, and 30% less about Topic 5, than other comparable white tracts (See Table 2).

While Topic 2 was rated as innocuous overall by participants in both the 'video' and 'video only' conditions, Figure 9A shows that post authors in all of these topics (2, 5, and 17) tend to frame posted videos as suspicious and criminal. The lack of significant regression results for other topics suggests that white enclaves versus white neighborhoods with nonwhite neighbors do not post significantly differently in other stranger-related topics.

5.2.8 Reports of Crime do not Correlate with Official Crime Rates. Ring is often marketed as a platform to report criminal activity, which is consistent with our finding that a major category that users post about is criminal activity, generally minor, such as package theft. However, we find minimal correlation between rates of posting about theft and official crime statistics: posting rates of Topic 1 (package theft) have no statistically significant correlation to theft or property/house theft; posting rates of Topic 3 (car break-ins) have no statistically significant correlation to theft, vehicle theft, or property/house theft; and posting rates of Topic 14 (stolen bikes and auto-theft) have no statistically significant correlation to property/house theft; and posting rates of Topics 4 (trespassing), 8 (breaking in) and 29 (burglary) have no major positive statistical correlation to any crime variables. In other words, the prevalence of content on Ring Neighbors reporting crimes does not reflect the official crime rates within that community.

5.2.9 The Impact of Owner-Occupancy. Unsurprisingly, owner-occupancy has a complex relationship to the kinds of posts on Ring Neighbors. In general, higher rates of owner-occupancy are related to an increase in posts that generally relate to homes, yards, and doors: all topics with a positive association with home ownership (2, 9, 13, 15, 17, and 18) consist of language such as "home", "door", "house", "neighbors", "knocking", or "driveway". Topic 13, which is characterized by words like "just", "guys", "neighbor", "alert", "safe", are generally community-focused posts with an alert or "keep safe" message. Owner-occupancy is the only demographic with a positive association with this topic, supporting the notion that owner-occupancy is related to an increase in community vigilance and identity.

5.2.10 Some Tracts Disproportionately Portray People of Color. There are some inconclusive patterns of racialized use on the platform. First, many predominantly white tracts appear to record people of color at disproportionate rates. Over 35% of posts made in tracts that are between 69% and 95% white (the first bar in Figure 12A) depict people of color. Our post codings also reveal some racial patterns. For instance, of the twelve most representative posts in Topic 2 where raters agreed on the race of people filmed in a video, 40% of them showed visibly non-white people, and the majority of these posts (4/6) were filmed in majority non-Hispanic white census tracts. Moreover, our regression shows that, for majority white tracts, a 10% increase in non-white neighbors increases Topic 2's prevalence by about 4%. However, we do not find any association between the percentage of posts that portray minority subjects and the racial makeup of a tract, including between non-white tracts and all kinds of white tracts.

6 DISCUSSION

By many measures, Amazon Ring is the fastest-growing corporate surveillance system in the US [9]. These porchfront doorbell cameras easily capture people and activity on nearby streets and sidewalks, making opting out of their gaze nearly impossible in areas with high adoption rates of Ring. For instance, it is impossible to walk around many neighborhoods in Los Angeles without being recorded by a Ring-affiliated device. Over 2,000 law enforcement agencies partner with Amazon Ring, giving these departments unprecedented access to surveillance videos and community networks [45]. There are major risks associated with such a widespread surveillance apparatus—especially one deeply intertwined with law enforcement—being a normalized part of everyday life: the increased risk of police use of facial surveillance technology in ways that disproportionately harm people of color and erode personal privacy; the "observers effect" it may incur on new generations of Americans growing up in neighborhoods that are constantly surveilled [12]; and the potential to perpetuate suburban neighborhood fears [20].

Yet, perhaps the most interesting thing about Ring is not just that these cameras exist, or that they are connected to policing networks. After all, it is not surprising that millions of consumers seek to surveil their own property. This is particularly true as retail shopping shifts from the storefront to consumers' doorsteps (the most common topic in Los Angeles is predictably, after all, about package theft). What makes Ring particularly unique is that it also encourages its users to treat this surveillance as a kind of *content*, and that it facilitates this relationship through its social network, Neighbors. Millions of hours of video content is recorded through Ring cameras every day, but only a small sliver is posted on the network. In this work, we try to measure some pieces of that network. We characterize where users post from and what they post, and use a case study of Los Angeles as a grounding example to explore how Neighbors is used in a major metro area in the US.

At both the national level and in Los Angeles, we find that an area's racial makeup is associated with posting rates on the Neighbors platform. In both cases, we find strong associations between whiteness and posting rates. In Los Angeles, we also see strong evidence that users use framing techniques to cast the videos they post of others as suspicious or criminal. These findings, further discussed below, provide what we see as inconclusive but suggestive evidence that one of Ring's main uses is as a kind of racialized gatekeeping tool. Patterns of criminalizing language and rates of claimed crime also suggest to us that Ring users blur the line between police work and their role as users and citizens.

6.1 The Impact of Race on Ring Use and Spread

We consistently find evidence both nationally and in Los Angeles that a neighborhood's racial makeup is related to how often its residents post, and what they post about. Generally, we find that whiter areas tend to use the platform more: on a national level, whiter counties post on Ring at a higher rate and, in our Los Angeles case study, white majority census tracts bordering only other white tracts (what we term "white enclaves") post on Ring at significantly higher rates than other tracts. In L.A., rates of property ownership are also positively correlated with Ring usage. This makes it clear that Ring is not used equally, but more often by a specific type of community: white, propertied enclaves.

We also identify some patterns in usage that might be interpreted as evidence of racial gatekeeping on the platform, but the evidence is not wholly conclusive. First, Ring posting rates are positively correlated with 311 calls to sweep homeless encampments in majority-white tracts in LA. Such calls bear the closest resemblance to the notion of neighborhood gatekeeping—they literally entail policing presence and belonging in a neighborhood. Second, we see some patterns from our STM regression that may support the gatekeeping hypothesis. White neighborhoods that share borders with non-white areas post more frequently in 3 of the 10 topics we identify as being about 'strangers'. Each of these topics receives a heavy dose of 'spin': videos are consistently rated less suspicious than videos with accompanying text.

These two results reveal a nuanced narrative. While white areas that border non-white neighborhoods post less than white 'enclaves', they post more about strangers in some topics where users heavily frame recorded activity as suspicious. Meanwhile, the tracts we identify as white enclaves post about community safety topics at much higher rates. If Neighbors is used as a racial gatekeeping tool by white tracts, this evidence is consistent. One would expect white tracts that border non-white areas to perform more of this gatekeeping than tracts in a racially homogenous enclave. We leave the question of why white 'enclave' tracts might post more about community safety concerns versus other topics to future work.

While consistent with the gatekeeping hypothesis, this evidence is not wholly conclusive. In other areas that we might expect to see evidence of this behavior, we find none. For example, we find no significant differences in the race of recorded subjects in different kinds of tracts. We also do not find strong patterns of racialized language in our topic analysis except for Topic 11, which shows no statistical relationship in our STM analysis. Additionally, although we see some disproportionate minority representation in post videos (See Section 5.2.10), we find no relationship between the race of video subjects and tract type (white 'enclave' vs white bordering tract).

That being said, we do come away from our deep reading with a strong sense that race is an undercurrent in how users frame their posts. Many posts we coded used euphemistic and racialized language to refer to subjects, such as "baggy clothes" or "dark". Although these patterns do not surface statistically, the content

is still there. For example, the eleventh most common topic in Los Angeles uses racialized language to describe "suspicious" people, many of whom are not filmed doing suspicious or explicitly criminal activities. In one striking example of this pattern, a video of a young adult on a scooter—whose race is unclear from the video but is nevertheless labeled Hispanic—is described as a "BREAK IN" (see Figure 8D and Appendix A #16).

6.2 Platform-Driven Paranoia and Framing

While videos are an important part of the Neighbors platform, the text that users attach to videos plays a crucial role. In our experimental survey, independent coders rated videos in almost all topics as more suspicious when they are also shown the accompanying post text. Posters use the text attached to videos to perform a kind of active framing that casts their video subjects as suspicious. This framing shows a kind of paranoia that pervades the Ring platform and that has been written about elsewhere [6].

We see this paranoia in major topics found in our modeling, as well. The second most common topic in Los Angeles frames the ordinary act of knocking on a door as suspicious. Many posts about "suspicious strangers" coded by raters and reviewed by the authors claim criminal activity is occurring without providing filmed evidence.

Does using Ring or the Neighbors platform make members of a community more likely to view mundane activities as threatening, criminal or suspicious? Unfortunately, our analysis does not allow us to determine whether users of Ring Neighbors are more likely than the general population to, for example, judge someone knocking on their door as out of place. This means that it could feasibly be the case that people who are inclined to use Neighbors are also simply more inclined to find a variety of innocuous behavior more threatening or suspect.

However, Ring is a massive, growing platform, and Neighbors is a primary way that users are incentivized to engage with both the company and the product. We do not find it far-fetched to speculate that exposure to Neighbors, and to the suspicious framing active on the platform, could shift a user's perspective towards the paranoid. To answer this question, future work could examine the extent to which users of the platform are more likely to view other activity as threatening, or the long-term impacts of using a platform where so much content is framed in terms of suspicion.

6.3 Participatory and Platformed Mass Surveillance

Ring can also be contextualized in a broader trend of "participatory mass surveillance", where people voluntarily surveil themselves and their neighbors in order to feel more secure. While most people in the United States would likely object to the government installing a nationwide system of cameras recording almost every street corner, consumers have ironically constructed this very network by purchasing products like Ring cameras. This phenomenon is not just limited to Ring's doorbell cameras as, in recent years, there has been a proliferation of increasingly invasive surveillance products. For instance, Ring has launched indoor cameras, the television service RingTV that shows content captured on Ring, and even a flying camera that "can see every angle in your home" [39]. Similarly, the hyper-local social network Citizen (formerly "Vigilante") piloted a live-streaming service to catch suspected criminals on air, which was notoriously used to conduct a manhunt of an innocent person in Los Angeles [14]. Other startups, such as Flock Safety, promise to "eliminate nonviolent crime" and sell cameras to homeowner's associations that automatically detect

when non-residents drive through a community [13]. Though these are just a few examples of participatory mass surveillance products, they reflect the growing trend of technologies that promise safety via policing and always-on surveillance and raise similar concerns to Ring: eroding privacy, fueling a culture of distrust in neighbors and viewing "outsiders" as threats or potential criminals. Examples such as RingTV or Citizen's live-streaming service also illustrate how surveillant content has even become a form of entertainment as well, reflecting a disturbing normalization of surveillance and policing in our everyday lives.

6.4 Blurring of Police Work and Citizen Surveillance

Our results also suggest that Ring functions as an extension of formal law enforcement, with Ring users taking on informal policing responsibilities. Police partnerships, in particular, enable police to directly request Ring footage from users without a warrant, expanding the amount of data law enforcement agencies have access to. In addition, the second most prevalent meta-category of content consists of reports of criminal activity, with the most commonly reported incidents being less severe crimes such as package theft. Importantly, we do not find any major, statistically significant correlation between posting rates of topics reporting criminal activity and official crime statistics. Reports of crime are certainly a reasonable use of a social network; however, the key distinction between a hyper-local platform such as Ring and other social networks is the knowledge that police in your community may directly act upon Ring posts. In this way, Ring users function as eyes and ears for local police departments, dramatically expanding the scope of policing. This is especially dangerous because Ring users often exaggerate the severity of suspiciousness or claim criminal evidence is occurring without definitive proof, though we leave it to future work to explore the relationship between content on platforms like Ring and policing responses in more detail.

In a small number of cases, Ring users will also call the police in response to activity they believe to be suspicious. A majority of posts in the 22nd most common topic on Ring contain references to calling the police, usually over minor or non-criminal activity. One notable example shows a person walking up the post author's driveway, filling up two water bottles at an exterior faucet, then leaving. The post is labeled "crime", is titled "man enters property and steals water," and the post author wrote that "911 was called" in the description (see Figure 8 B and Appendix A #6). While this incident is technically a crime, we believe it reflects a tendency on Ring to film minor incidents, upload them online, and call the police.

6.5 Possible Interventions and Design Lessons

Unsurprisingly, we also find that Ring is used as a community messaging board to discuss topics such as gunshot sounds, lost pets, or police in the vicinity. This usage of Ring is consistent with how the platform markets itself as a neighborhood messaging board, though posts about "suspicious" strangers or crime appear more frequently.

Ultimately, though our results suggest that Ring can certainly be used for productive purposes such as informing the community, there are significant risks associated with the platform. In particular, we believe that Ring can perpetuate a culture of paranoia by priming Ring users to see neighbors as threats. We find it likely that the effect of user framing on posts and the psychological effects of frequent reports about criminal activity—which we have shown is not correlated with official crime statistics—distorts peoples' perception of their community, though we leave it to future research to causally answer this question.

Similarly, though we do not test this question empirically, we come away from our grounding reading with the impression that many posts on Ring show activity beyond the limits of a person's property such as drivers or people passing by on sidewalks. While each individual post may seem minor, we worry that the spread of Ring normalizes the surveillance of public areas, which in turn can shift legal standards around reasonable expectations of privacy.

Our results also suggest possible interventions to improve a platform like Ring. We find that framing—the way posters describe activity—is important in shaping how other people perceive Ring posts, and by extension, their community. We also speculate that a culture of paranoia and distrust on Ring, often combined with racial biases, further distorts how Ring users view their neighbors. Therefore, one potential countermeasure is for the platform to prompt more accurate framing: for instance, when Ring users draft posts, they could be reminded to be as accurate, unbiased, and respectful of any filmed subject's privacy as possible. Existing research has found that prompting social media users to consider accuracy before sharing a post can improve the quality of the news they share [36]. Likewise, the hope is that by reminding Ring users of normative value, framing bias on the platform can be reduced. Additionally, our research suggests that users often share footage of innocuous activity. There are harms associated with this practice: people such as solicitors knocking on doorbells, teenagers riding their bikes outside, or Amazon delivery workers are all at risk of being filmed without consent, and it would be helpful for Ring users to also consider these effects when uploading posts. Adding friction, like the prompts described previously, could also help Ring users consider whether they actually want to upload certain videos.

There are many other aspects of Ring, as a product and surveillance network, that go unexamined in this work. As a product that is often marketed as a way to deter package theft, Ring can also be framed as a way that online retailers like Amazon translate retail surveillance practices onto residential doorsteps. Although there are few videos that depict Amazon and package delivery workers, they are no doubt some of the most-recorded people on Ring as a platform. Future work might include them in a more detailed analysis. This study is also limited as an observational analysis and case study. There are causal questions about Ring as a platform that could be answered with other designs. For example, how does continued use of a platform like Ring impact residents' perception of community safety?

6.6 Data Use and Ethics

A potential limitation of using Neighbors posts as an instrument of analysis are the inherent biases within the dataset. Because the dataset was collected by scraping posts from Ring, it is not representative of the entire population of people who use Ring. Though Ring does not release statistics on what proportion of users actually post to Neighbors, it is likely that people who post on Neighbors behave differently from other Ring users. As discussed previously, our finding that Ring posters tend to see innocuous activity as suspicious could also be explained by the sampling bias of our dataset: those who purchase security devices, regularly monitor them, and decide to post footage are more likely to be suspicious of events around their property. This is still a relevant finding, especially given that factors such as race and property ownership correlate with Ring usage, but is important to contextualize within how the dataset is produced. Moreover, the dataset is not representative of all activity filmed on Ring cameras, only the activity that users choose to upload.

There are also numerous potential ethical concerns associated with the use of Ring Neighbors data. Researchers such as Buck et al. have argued that using "public" data, such as that on social networks, can violate people's reasonable expectation of privacy, pose challenges related to confidentiality and anonymity, and draw unwanted attention to vulnerable demographics [10]. In considering these risks, we have taken steps to protect participant anonymity and believe that our use of Ring data is consistent with a Neighbors user's reasonable expectation of privacy. First, Ring Neighbors posts are not full-text searchable and posts are not attached to a user profile on the Neighbors platform. Additionally, we are not publicly releasing our dataset of Neighbors posts, and for any post mentioned in the paper, we have removed potentially identifying features such as exact location or faces in photos. This ensures that any post cited in our paper cannot be linked to a real person. Second, because we de-identify the data presented, we also believe that we are not exposing individuals to unwanted attention or harm. Lastly, we also believe that our use of the data is consistent with a Ring user's reasonable expectation of privacy. Ring users have an assumption, as outlined in the terms of service, that their data and content may be used by law enforcement, government officials, and other third parties. While users may not expect their data to be used by researchers, we argue that there is a public benefit in performing this analysis and the use of this data is justifiable if kept anonymized.

7 CONTRIBUTIONS

In this paper, we characterize the spread of Amazon Ring, arguably the fastest-growing private surveillance network in America. We provide a critical summary of Ring, situating it within long-standing conversations on surveillance and community self-policing, and characterize it as a form of "participatory mass surveillance". We then use a data-driven approach at the national level and in a case study of Los Angeles to investigate which communities are more likely to use the platform. We find in both cases that whiter neighborhoods post more on Neighbors, with white 'enclaves' in Los Angeles posting far more than other areas. Using a topic modeling approach and an experimental survey design, we analyze the most "representative" posts from our topics in Los Angeles, showing that users often elevate otherwise innocuous or minor behavior to levels of suspicion and criminality. We also find that white neighborhoods that border non-white areas post more in some topics about strangers, supporting the idea that Neighbors is used as a racial gatekeeping tool.

Ring advertises itself as a community safety tool and a way for neighbors to communicate important information to one another, and the platform is certainly used this way by some. Yet we overwhelmingly see in our analysis of Neighbors content and in our spatial modeling that Ring content can often be racialized, and is used to report and amplify highly subjective accounts of suspicious behavior. As platforms like Ring continue to integrate into communities across the US, we hope this work can provide a grounded context for informed discussions about Ring, the Neighbors platform, and other technologies like it.

REFERENCES

[1] [n. d.]. 2020 Greater Los Angeles Homeless Count. ([n. d.]). https://www.lahsa.org/documents?id=4558-2020-greater-los-angeles-homeless-count-presentation
[2] [n. d.]. FBI Uniform Crime Reporting. https://www.fbi.gov/services/cjis/ucr/
[3] [n. d.]. Introducing the Neighbors App: The New Neighborhood Watch. https://blog.ring.com/products-innovation/introducing-the-neighbors-app-the-new-neighborhood-watch/
[4] [n. d.]. Neighbors App by Ring. https://ring.com/neighbors
[5] [n. d.]. QuickFacts: Los Angeles City, California. https://www.census.gov/quickfacts/losangelescitycalifornia

[6] [n. d.]. Ring Camera Surveillance Is Transforming Suburban Life. *Wired* ([n. d.]). https://www.wired.com/story/ring-surveillance-suburbs/

[7] Mark Andrejevic. [n. d.]. The Work of Watching One Another: Lateral Surveillance, Risk, and Governance. 2, 4 ([n. d.]).

[8] Sam Biddle and 6:00 A.m. [n. d.]. LAPD Sought Ring Home Security Video Related to Black Lives Matter Protests. ([n. d.]). https://theintercept.com/2021/02/16/lapd-ring-surveillance-black-lives-matter-protests/

[9] Lauren Bridges. [n. d.]. Infrastructural Obfuscation: Unpacking the Carceral Logics of the Ring Surveillant Assemblage. 0, 0 ([n. d.]), 1–20. https://doi.org/10.1080/1369118X.2021.1909097

[10] Amber M. Buck and Devon F. Ralston. 2021. I didn't sign up for your research study: The ethics of using "public" data. *Computers and Composition* 61 (09 2021), 102655. https://doi.org/10.1016/j.compcom.2021.102655

[11] Dell Cameron. [n. d.]. Everything Cops Say About Amazon's Ring Is Scripted or Approved by Ring. ([n. d.]). https://gizmodo.com/everything-cops-say-about-amazons-ring-is-scripted-or-a-1836812538

[12] Dell Cameron and Dhruv Mehrotra. [n. d.]. Ring's Hidden Data Let Us Map Amazon's Sprawling Home Surveillance Network. ([n. d.]). https://gizmodo.com/ring-s-hidden-data-let-us-map-amazons-sprawling-home-su-1840312279

[13] Joseph Cox. [n. d.]. Inside TALON. https://www.vice.com/en/article/bvx4bq/talon-flock-safety-cameras-police-license-plate-reader

[14] Joseph Cox and Jason Koebler. 2021. 'FIND THIS FUCK:' Inside Citizens Dangerous Effort to Cash In On Vigilantism. https://www.vice.com/en/article/y3dpyw/inside-crime-app-citizen-vigilante

[15] Mike Davis. [n. d.]. *City of Quartz: Excavating the Future in Los Angeles*. Verso.

[16] Sheena Erete and Jennifer O. Burrell. [n. d.]. Empowered Participation: How Citizens Use Technology in Local Governance. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (2017). 2307–2319.

[17] Chris Gilliard. [n. d.]. Caught in the Spotlight. https://urbanomnibus.net/2020/01/caught-in-the-spotlight/

[18] Matthew Guariglia and Dave Maass. [n. d.]. LAPD Requested Ring Footage of Black Lives Matter Protests. https://www.eff.org/deeplinks/2021/02/lapd-requested-ring-footage-black-lives-matter-protests

[19] Caroline Haskins. [n. d.]. Amazon's Home Security Company Is Turning Everyone Into Cops. ([n. d.]). https://www.vice.com/en/article/qvyvzd/amazons-home-security-company-is-turning-everyone-into-cops

[20] Caroline Haskins. [n. d.]. How Ring Transmits Fear to American Suburbs. ([n. d.]). https://www.vice.com/en/article/ywaa57/how-ring-transmits-fear-to-american-suburbs

[21] Peter Holley. [n. d.]. This Patent Shows Amazon May Seek to Create a 'Database of Suspicious Persons' Using Facial-Recognition Technology. ([n. d.]). https://www.washingtonpost.com/technology/2018/12/13/this-patent-shows-amazon-may-seek-create-database-suspicious-persons-using-facial-recognition-technology/

[22] University of Pennsylvania Jacob Kaplan. [n. d.]. Jacob Kaplan's Concatenated Files: Uniform Crime Reporting Program Data: Offenses Known and Clearances by Arrest, 1960-2019. https://doi.org/10.3886/E100707V16

[23] Rahim Kurwa. [n. d.]. Building the Digitally Gated Community: The Case of Nextdoor. 17, 1/2 ([n. d.]), 111–117. https://doi.org/10.24908/ss.v17i1/2.12927

[24] Victoria Law and Maya Schenwar. [n. d.]. *Prison by Any Other Name: The Harmful Consequences of Popular Reforms*. The New Press.

[25] Joscha Legewie and Merlin Schaeffer. [n. d.]. Contested Boundaries: Explaining Where Ethnoracial Diversity Provokes Neighborhood Conflict. 122, 1 ([n. d.]), 125–161.

[26] Amy E. Lerman and Vesla Weaver. [n. d.]. Staying out of Sight? Concentrated Policing and Local Political Action. 651, 1 ([n. d.]), 202–219. https://doi.org/10.1177/0002716213503085

[27] Karen Levy and Solon Barocas. [n. d.]. Privacy at the Margins/Refractive Surveillance: Monitoring Customers to Manage Workers. 12 ([n. d.]), 23.

[28] Dan A. Lewis and Michael G. Maxfield. [n. d.]. Fear in the Neighborhoods: An Investigation of the Impact of Crime. 17, 2 ([n. d.]), 160–189.

[29] Louise Matsakis. [n. d.]. Ring Camera Surveillance Is Transforming Suburban Life. ([n. d.]). https://www.wired.com/story/ring-surveillance-suburbs/

[30] David Mimno, Hanna Wallach, Edmund Talley, Miriam Leenders, and Andrew McCallum. [n. d.]. Optimizing Semantic Coherence in Topic Models. In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing* (2011). 262–272. https://dirichlet.net/pdf/mimno11optimizing.pdf

[31] Anouk Mols and Jason Pridmore. [n. d.]. When Citizens Are "Actually Doing Police Work": The Blurring of Boundaries in WhatsApp Neighbourhood Crime Prevention Groups in The Netherlands. 17, 3/4 ([n. d.]), 272–287. https://doi.org/10.24908/ss.v17i3/4.8664

[32] Gaia Mosconi, Matthias Korn, Christian Reuter, Peter Tolmie, Maurizio Teli, and Volkmar Pipek. [n. d.]. From Facebook to the Neighbourhood: Infrastructuring of Hybrid Community Engagement. In *Computer Supported Cooperative Work (CSCW)* (2017), Vol. 26. Springer, 959–1003.

[33] Laura K. Nelson. [n. d.]. Computational Grounded Theory: A Methodological Framework. 49, 1 ([n. d.]), 3–42. https://doi.org/10.1177/0049124117729703
[34] Daniel O'Brien. [n. d.]. *The Urban Commons — Daniel T. O'Brien*. Harvard University Press. https://www.hup.harvard.edu/catalog.php?isbn=9780674975293
[35] Will Payne. [n. d.]. Welcome to the Polygon: Contested Digital Neighborhoods and Spatialized Segregation on Nextdoor. 6 ([n. d.]). http://computationalculture.net/welcome-to-the-polygon-contested-digital-neighborhoods-and-spatialized-segregation-on-nextdoor/
[36] Gordon Pennycook, Ziv Epstein, Mohsen Mosleh, Antonio A. Arechar, Dean Eckles, and David G. Rand. 2021. Shifting attention to accuracy can reduce misinformation online. *Nature* 592, 7855 (04 2021), 590–595. https://doi.org/10.1038/s41586-021-03344-2 Number: 7855 Publisher: Nature Publishing Group.
[37] Joel R Reidenberg. [n. d.]. Privacy in Public. ([n. d.]), 21.
[38] Emily Alpert Reyes, Benjamin Oreskes, and Doug Smith. [n. d.]. L.A. Is Swamped with 311 Complaints over Homeless Camps. But Are the Cleanups Pointless? ([n. d.]). https://www.latimes.com/local/lanow/la-me-ln-homeless-cleanups-accelerate-20190607-htmlstory.html
[39] Ring. 2021. Ring Always Home Cam. https://www.youtube.com/watch?v=Kk4HOGa257k
[40] George Ritzer and Nathan Jurgenson. [n. d.]. Production, Consumption, Prosumption: The Nature of Capitalism in the Age of the Digital 'Prosumer'. 10, 1 ([n. d.]), 13–36.
[41] Margaret Roberts, Brandon Stewart, and Dustin Tingley. [n. d.]. Stm: An R Package for Structural Topic Models. 91, 2 ([n. d.]), 1–40. https://doi.org/10.18637/jss.v091.i02
[42] Jonathan Robinson, Cheskie Rosenzweig, Aaron J. Moss, and Leib Litman. 2019. Tapped out or barely tapped? Recommendations for how to harness the vast and largely unused potential of the Mechanical Turk participant pool. *PLOS ONE* 14, 12 (12 2019), e0226394. https://doi.org/10.1371/journal.pone.0226394 Publisher: Public Library of Science.
[43] Ben Fox Rubin. [n. d.]. How Ring's Neighbors App Is Making Home Security a Social Thing. https://www.cnet.com/home/security/how-rings-neighbors-app-is-making-home-security-a-social-thing/
[44] Alexandra Schofield, M\ aans Magnusson, Laure Thompson, and David Mimno. [n. d.]. Understanding Text Pre-Processing for Latent Dirichlet Allocation. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics* (2017), Vol. 2. 432–436.
[45] Jeffrey Shen. [n. d.]. Amazon Ring's Police Partnership Tracker. https://ring.jeffreyshen.com/
[46] Keith Spiller and Xavier L'Hoiry. [n. d.]. Watch Groups, Surveillance, and Doing It for Themselves. 17, 3/4 ([n. d.]), 288–304. https://doi.org/10.24908/ss.v17i3/4.8637
[47] Jay Stanley. [n. d.]. Should You Buy a Ring Doorbell Camera? https://www.aclu.org/blog/privacy-technology/surveillance-technologies/should-you-buy-ring-doorbell-camera
[48] Luke Stark and Karen Levy. [n. d.]. The Surveillant Consumer. 40, 8 ([n. d.]), 1202–1220. https://doi.org/10.1177/0163443718781985
[49] Sarita Yardi and Danah Boyd. [n. d.]. Dynamic Debates: An Analysis of Group Polarization over Time on Twitter. 30, 5 ([n. d.]), 316–327.
[50] Bilge Yesil. [n. d.]. Watching Ourselves: Video Surveillance, Urban Space and Self-Responsibilization. 20, 4-5 ([n. d.]), 400–416. https://doi.org/10.1080/09502380600708770
[51] Shoshana Zuboff. [n. d.]. Big Other: Surveillance Capitalism and the Prospects of an Information Civilization. 30, 1 ([n. d.]), 75–89.

APPENDIX

Table 3. Appendix A: Referenced Neighbors Posts

ID	Topic	Gamma	Title	Description	Category
1	2	0.5667075	Shady guy knocking instead of ringing doorbell	Who is this guy and why does he knock on the door instead of ringing the doorbell ? Seems shady to me.	stranger
2	2	0.5734577	I don't expect anybody !	This stranger man ring on my s stranger man ring on my door bell	stranger
3	2	0.5648080	Has he came to anyone elses house?	He came twice today knocking tried speaking with him through the mic but he wouldn't respond. Couldn't tell if he was from some company or who he was.	stranger
4	2	0.5220549	Anyone recognize this giy	Wasn't home and couldn't tell from my Ring if he left (I guess I have to record longer) but wondering if he came to anyone else's door. Is he just an unwanted solicitor?	stranger
5	22	0.6107924	3 Young Men Fight w/ Bricks	2 young men were teaming up on another young man who had a brick trying to defend himself. Tried calling the cops 4 times & got a busy signal. What a joke	crime
6	22	0.5421076	man enters property and steals water	This man entered the property and begins filling up water bottles using our exterior driveway faucet. After the alarm was triggered he headed east on Overland apparently Pushing a shopping cart. 911 was called	crime

Table 3. Appendix A: Referenced Neighbors Posts *(continued)*

ID	Topic	Gamma	Title	Description	Category
7	22	0.5039263	Teenagers slamming violently into front doors	There's a large group of teenagers roaming Torch and Valerie area slamming hated into doors, threatening people, lighting off firecrackers, etc. police have been called by us and at least one other neighbor.	stranger
8	22	0.6152284	Trespassers using my water	At 312 last night I was woken up by these two men using my front water spout to "wash their hair" as they "just dyed it and needed to rinse it out". I told them to get off my property immediately and he kept saying just another minute! I told him he's welcome to stay as long as he would like then and explain himself to the police as they were on the way. I hadn't actually called the police because I had a feeling I could scare them off on my own and I did. Thankful for my ring alerting me to these two intruders!	suspicious

Table 3. Appendix A: Referenced Neighbors Posts *(continued)*

ID	Topic	Gamma	Title	Description	Category
9	22	0.4710522	Walking down our neighborhood which is a dead end	I witnessed this guy walking down my street which is a dead end There was a amazon truck making deliveries. I rolled up to this piece of work and asked if I could help him. He told me to mind my business. Said if he had a gun in his backpack he could shoot me. I told him that would be a bad move. Called the police and talked to them and guess what. NOTHING. WERE on our own guys ! Call the wla police station and ask for the chief Complain Call strangers out Don't be silent. We need to take back this neighborhood. I'll be the front man. I need support !!	suspicious
10	5	0.4576365	Suspicious character on Sawtelle	This suspicious character looks like he's up to no good. Possibly scoping out condos while owners are away enjoying the holidays. Keep an eye out for this guy!	suspicious
11	5	0.3498293	Suspicious guy walking in Roscomare Rd	Suspicious guy walking up and down Roscomare Rd. Has a car parked on the street. Looks like scoping out houses.	suspicious
12	5	0.4463386	woman on my porch at 2:51 am	Does she look familiar? She looks like she was looking for something to steal.	suspicious
13	17	0.3058885	No shame	Probably saw it was from kohl's and ditched it [emoji][emoji]	suspicious

Table 3. Appendix A: Referenced Neighbors Posts *(continued)*

ID	Topic	Gamma	Title	Description	Category
14	17	0.3887667	Suspicious person	At 1:30 am. I was walking home from work and a nicely dressed young black man walking in the opposite direction as myself. Asked me if I had a phone charger. I told him no, sorry. After about half of a block I hear miss, miss. He was following me. I pulled up to the nearest doorway. It was a small hotel. He continued to follow me saying miss. I reached into my bag. Like I was looking for my keys. And grabbed my tazer. He asked me if I could charge his phone. That he need to call an uber. I told him sorry bit my husband wouldn't like that. He dissapired around some bushes. I acted like I was talking my phone. Trying to keep an eye on him. After a few minutes I carefully walked home. He was no where to be seen. Not in any direction. This was by the magic castle. People please watch your back when walking in this area.	suspicious

Table 3. Appendix A: Referenced Neighbors Posts *(continued)*

ID	Topic	Gamma	Title	Description	Category
15	11	0.7541415	Another construction thief	Description of the person: African American male Adult age Approximately 5'9" tall Dark complexion He's wearing: Black backpack, medium Black beanie with no logos Black jacket with white piping on the sleeves Grey/blue baggy pants, sagged like a thug Red high tops with red laces and white bottoms He: Smokes cigarettes, either light or menthol Large gold ring on left ring finger Smart enough to wear an orange vest Stores his phone in his right hip pocket Let's the white earphones hang out of his pocket Went to the 77th to make a report, but as the trucks weren't mine they wouldn't take a report.	suspicious
16	11	0.6579360	BREAK IN	7/30/2018 3:30pm Male Hispanic early 20's on a SCOOTER Dark shirt, light colored shorts, white shoes and black socks. Please contact me or police if seen or known.	crime

Table 4. Appendix B: Description of Spatial Regression Variables

Variable Name	Description
actual mtr veh theft total	# of incidents of motor vehicle theft
actual robbery total	# of incidents of robbery
actual theft total	# of incidents of theft
actual assault total	# of incidents of assault
actual burg total	# of incidents of burglary
non hispanic white	% of population that is non-Hispanic white
black	% of population that is Black/African American
hispanic or latino	% of population that is Hispanic/Latino
median household income	Median household income
aian alone	% of population that is American Indian/Alaska Native
asian alone	% of population that is Asian
two or more races	% of population that is two or more races
nhpi alone	% of population that is Native Hawaiian/Pacific Islander
some other alone	% of population that is some other race
owner occupied	% of occupied housing units that are owner-occupied
total occupied	# of occupied housing units
median property value	Median property value

Table 5. Appendix C: Top 30 Topics in Structured Topic Model

Topic	Gamma	Label	Description	Top 10 Words
1	0.0438760	Package theft	Posters reporting their packages being stolen, often with video of the theft.	package, stole, porch, packages, mail, front, stealing, thief, amazon, box
2	0.0432110	Stranger knocking on door	Strangers knocking on doors. Authors often imply strangers are pretending to be workers. Most people recorded in top 20 posts are non-white.	door, doorbell, rang, came, bell, knocked, man, strange, answer, home
3	0.0410743	Car break-ins and theft	Most posts reviewed describe car break-ins.	car, night, last, cars, morning, break, broken, anything, locked, open
4	0.0374387	Stranger looking around or trying to enter	Strangers looking around the author's property, often trying to enter the house or backyard. Authors often tell their neighbors to beware or be on the lookout.	guy, around, tried, trying, recognize, home, today, lookout, beware, came
5	0.0355432	Suspicious person who wants to commit a crime	Posts and usually videos of people being "suspicious." Posts do not usually contain explicit criminal activity, but authors tend to ascribe criminal intent to the people being recorded.	looking, person, like, suspicious, looks, open, seems, houses, taking, walks
6	0.0329566	Asking questions	Questions about neighborhood activity, spanning a variety of topics such as police, strangers, and lost cats.	anyone, know, else, don't, let, anything, someone, knows, thanks, wondering
7	0.0279756	Sharing videos	Posts about suspicious or criminal activity, with most posts mentioning phrases related to sharing video footage or pictures. Many authors also report incidents and share footage with the police.	video, time, report, another, neighbor, second, police, face, first, post

Table 5. Appendix C: Top 30 Topics in Structured Topic Model *(continued)*

Topic	Gamma	Label	Description	Top 10 Words
8	0.0274530	Climbing over fences	Strangers breaking into backyards, often by climbing over fences and breaking through gates.	back, came, went, yard, gate, backyard, neighbors, fence, jumped, side
9	0.0268634	Person or car outside house	Strangers or cars lingering outside of people's homes	house, front, window, looked, side, right, door, outside, look, mins
10	0.0266144	Describing security systems	Posts cover a range of topics, but most posts mention some aspect of the author's security system, for instance, describing how motion lights went off or Ring failed to capture an event.	ring, camera, didn't, caught, cameras, also, didnt, security, unfortunately, motion
11	0.0245495	Suspicious people described with racialized language	Posts mostly describe activity that the authors perceive as suspicious or criminal. Almost all posts reviewed explicitly mention race, and many use racialized language.	black, male, wearing, white, shirt, backpack, blue, dark, hair, hoodie
12	0.0245351	Missing pets or relatives	Authors asking for help finding lost pets or, less commonly, missing people. Language used is significantly politer, and phone numbers are often included. Most posts in top 20 posts do not include videos. Several police notifications are also included in this category.	please, help, call, thank, information, dog, missing, find, contact, lost
13	0.0239729	Various, but community-focused	Community focused posts that read like authors are posting to a small community.	just, guys, neighbors, alert, safe, everyone, aware, stay, thought, wanted

Table 5. Appendix C: Top 30 Topics in Structured Topic Model *(continued)*

Topic	Gamma	Label	Description	Top 10 Words
14	0.0229948	Bike and car theft	Authors reporting theft of their own property, with most posts being about car or bicycle theft, with some package theft.	stolen, stole, broke, someone, bike, got, garage, stuff, two, things
15	0.0223978	Neighborhood commotion	Describes commotion in the area, such as blockages, smoke, or lots of police. Posts are usually framed as questions. Most posts in top 20 posts do not include videos.	going, around, anybody, cars, whats, cops, police, idea, theres, blocked
16	0.0222927	Loud sounds	Describes gunshots, fireworks, explosions, or other loud sounds. Posts are usually framed as questions. Most posts in top 20 posts do not include videos.	hear, like, heard, anyone, loud, else, gunshots, just, sounded, fireworks
17	0.0210785	Suspicious walking-based activity	Posts mostly describe activity that the authors perceive as suspicious or criminal. Most posts reviewed contain phrases related to walking.	saw, walked, away, walking, driveway, turned, probably, started, soon, towards
18	0.0209920	Suspicious solicitor	Strangers knocking on authors' doors, often teenagers/young adults asking for money. Authors tend to imply the solicitors are scammers.	asked, said, asking, knocking, told, saying, didn't, selling, kid, name

Table 5. Appendix C: Top 30 Topics in Structured Topic Model *(continued)*

Topic	Gamma	Label	Description	Top 10 Words
19	0.0188974	Police activity and reports	Most posts reviewed were about police activity in area such as helicopters circling and shut-down streets. Some posts by police departments asking for help with cases. Many posts mentioned specific locations.	area, police, helicopter, activity, suspect, lapd, ventura, closed, reseda, individuals
20	0.0185304	Various iterations of phrase "can you see"	Posts cover a variety of unrelated topics. Most posts use phrases of the format "can you see" or "as you can see", topic likely grouped by linguistic patterns.	see, can, checking, check, someone, try, can't, think, tell, clearly
21	0.0179370	Various safety topics with street intersections	Posts cover a variety of safety related topics such as gunshots or police nearby. All posts reviewed contain street or intersection names, topic likely grouped by this linguistic pattern.	near, ave, blvd, shots, west, hills, school, valley, san, east
22	0.0175740	Calling 911	Mostly petty crime or suspicious people. Nearly all posts emphasize calling 911 or the police.	police, called, told, said, call, cops, leave, man, notified, trespassing
23	0.0172600	Suspicious/criminal activity, using word "seen"	Posts describe suspicious people or criminal activity such as alleged stalkers, package theft, and trespassing. Most posts reviewed used the word "seen", often as a question (i.e. has anybody seen this person?)	seen, neighborhood, never, 've, times, man, doesn't, knocks, home, late

Table 5. Appendix C: Top 30 Topics in Structured Topic Model *(continued)*

Topic	Gamma	Label	Description	Top 10 Words
24	0.0167308	Cars	Most posts mentions cars, but context varies. Mostly car related crime, suspicious people in cars, and specific car models.	car, parked, white, driving, van, drove, pulled, stopped, toyota, noticed
25	0.0166342	Solicitors and crime	Describes suspicious solicitors (often implied to be scammers) and crime such as break-ins. Many posts reviewed use tough-on-crime rhetoric and claim crime is on the rise.	people, come, home, homes, day, neighborhood, don't, need, third, crime
26	0.0154898	Various suspicious people	Posts mostly describe "suspicious" people. Unclear what about this topic is unique.	street, woman, took, picture, across, waiting, seemed, middle, sitting, min
27	0.0154874	Helicopters	Authors asking about helicopters in area.	now, helicopters, lot, circling, happening, right, lots, flying, hour, sirens
28	0.0153532	Animals	Posts are about a variety of animals: for instance, lost animals, dangerous wildlife, and aggressive pets.	coyote, cat, dog, running, around, saw, pets, animal, yard, careful
29	0.0151188	Burglary	Burglaries. Posts describe method of entry, such as windows, smashed, or screens. Many posts mention the items stolen, primarily jewelry but also money and electronics.	door, home, window, items, took, friday, taken, glass, ransacked, broke
30	0.0148504	Theft and vandalism, Spanish and English	Many posts are in Spanish. Describes crime, mostly theft and vandalism of the author's property.	happened, ago, nothing, days, taken, weeks, couple, month, new, months

Table 6. Appendix C: Significant (>95% level) STM regression results. Empty cells indicate no significant effect. Crime reporting rates are shown in incidences per 1,000 people. Property value is in tens of thousands of dollars. Percent nonwhite neighbors and owner occupation rates are in 1% units. For example, the result for owner occupied rates on Topic 2 can be read as: 'A 1% increase in owner occupancy rates is associated with a .43% increase in Topic 2 posting rates.'

topic	violent crime	vehicle theft	owner occupied	(maj. white) pct nonwhite neighbors	gun violence	(maj. white) w/ all white neighbors	median property value	theft	pct nonwhite neighbors	property or house theft
1	-0.0029% (0.00084%)***	0.0023% (0.00082%)**								
2			0.43% (0.13%)***	0.4% (0.2%)*						
3		0.00096% (0.00049%)*								
4	-0.00055% (0.00028%)*	0.00068% (0.00024%)**			-0.01% (0.0049%)*					
5		0.002% (0.00045%)***			-0.018% (0.0066%)**	-30% (12%)*	0.0000048% (0.0000024%)*	-0.0015% (0.00067%)*		
6		-0.0016% (0.0004%)***								
7	-0.0011% (0.00052%)*								-0.44% (0.16%)**	
8			-0.39% (0.12%)**		-0.027% (0.0089%)**			-0.0021% (0.00087%)*		
9	-0.0014% (0.00037%)***	0.001% (0.00035%)**	0.4% (0.087%)***							
12		-0.0018% (0.00079%)*			0.036% (0.016%)*			0.0031% (0.0015%)*		

Table 6. Appendix C: Significant (>95% level) STM regression results. Empty cells indicate no significant effect. Crime reporting rates are shown in incidences per 1,000 people. Property value is in tens of thousands of dollars. Percent nonwhite neighbors and owner occupation rates are in 1% units. For example, the result for owner occupied rates on Topic 2 can be read as: 'A 1% increase in owner occupancy rates is associated with a .43% increase in Topic 2 posting rates.' *(continued)*

13			0.39% (0.11%)***		-0.019% (0.0076%)*			
14		0.0018% (0.0007%)**	-0.74% (0.18%)***		-0.027% (0.012%)*		-0.0031% (0.0012%)*	
15	0.0032% (0.00097%)**					54% (24%)*	-0.000028% (0.0000096%)**	-0.0036% (0.0016%)*
16								-0.0036% (0.0017%)*
17	0.0011% (0.00051%)*		0.45% (0.14%)**	0.53% (0.22%)*		-43% (15%)**	0.000006% (0.000003%)*	-0.002% (0.00074%)**
18			0.9% (0.25%)***					

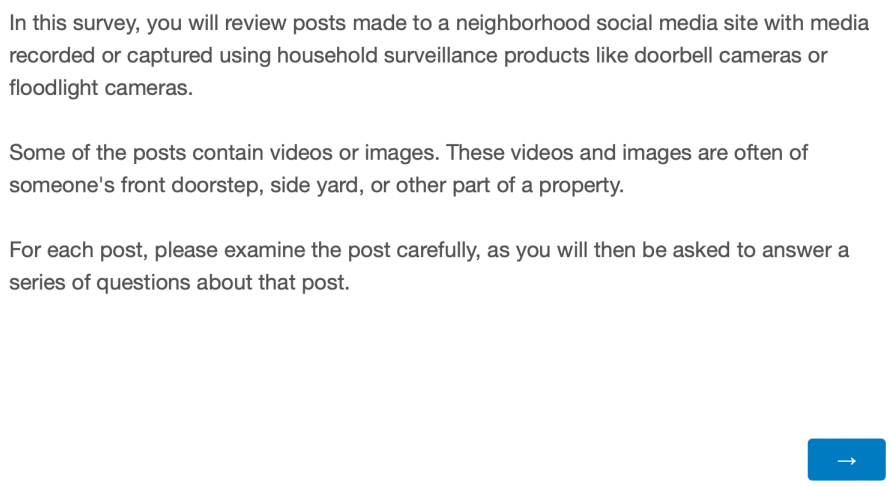

Fig. 13. Introduction to the survey

Appendix D: Survey Design

Our experimental survey was designed in Qualtrics. The figures below show example surveys a participant might have recieved under our two randomized conditions. In the first, participants saw both post text and video. In the second, participants saw only videos. Posts were randomized equally across all participants and conditions.

On Mechanical Turk, we calculate the average time for a single HIT by timing our own performance, and set pay so on average, workers make an equivalent of $15/hr.

Introduction.

Attention Checks. We ask participants to complete several attention checks. Any responses from participants that do not pass our two attention checks are not included in our analysis, but we still pay all participants.The first attention check asks participants to simply enter a number (Figure 15). The second asked participants to ignore the question text and select two specific answers from a list (Figure 16)

Post. After solving each attention check, participants were assigned to a condition randomly, and shown a series of posts. Figure 17 shows what a participant in the "video only" condition would see, including the "suspiciousness" likert question, while Figure 18 shows what a participant in the "video and text" condition would see.

Questions. After being shown the post, users would scroll down and answer the following series of questions. If a participant answered that a crime was shown or a crime was claimed by a poster, they would then be asked to rate the severity of that crime.

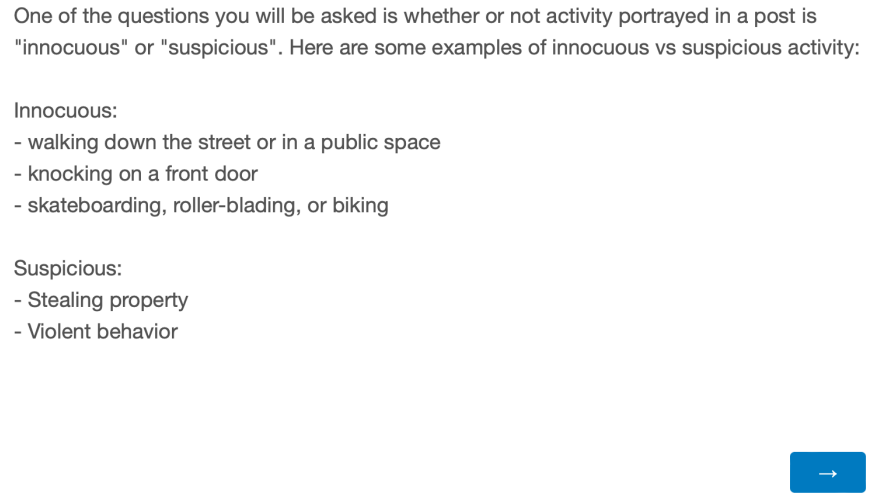

Fig. 14. Introduction to the survey, explaining our terms 'innocuous' and 'suspicious'

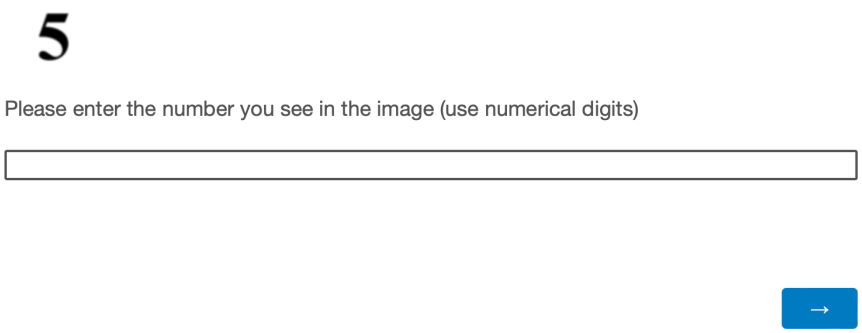

Fig. 15. Attention check 1

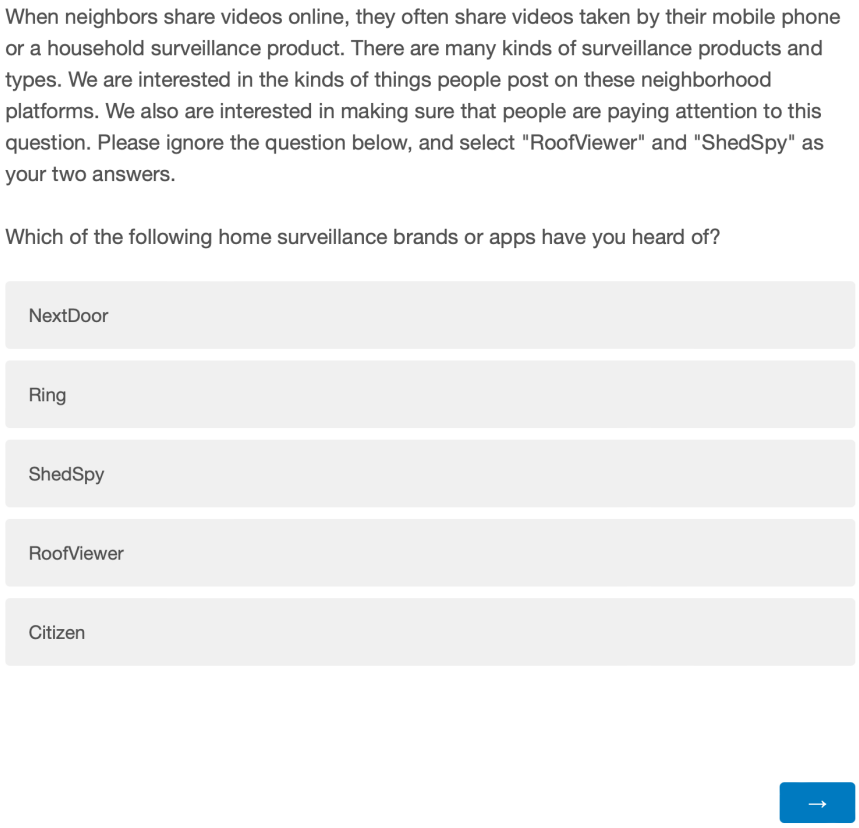

Fig. 16. Attention check 2

On a scale of 'Very Innocuous' to 'Very Suspicious', please rate how suspicious the activity portrayed in the post was.

- Very Innocuous
- Innocuous
- Somewhat Innocuous
- Neutral
- Somewhat Suspicious
- Suspicious
- Very Suspicious

Fig. 17. Video-only survey post.

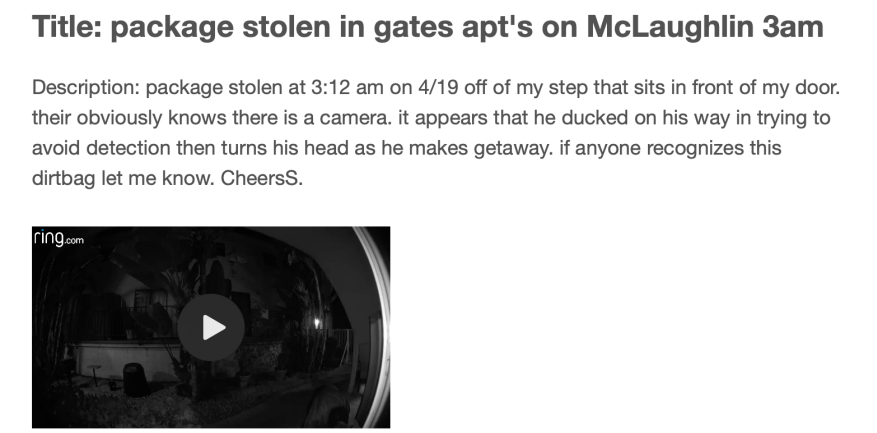

Fig. 18. Video and text condition.

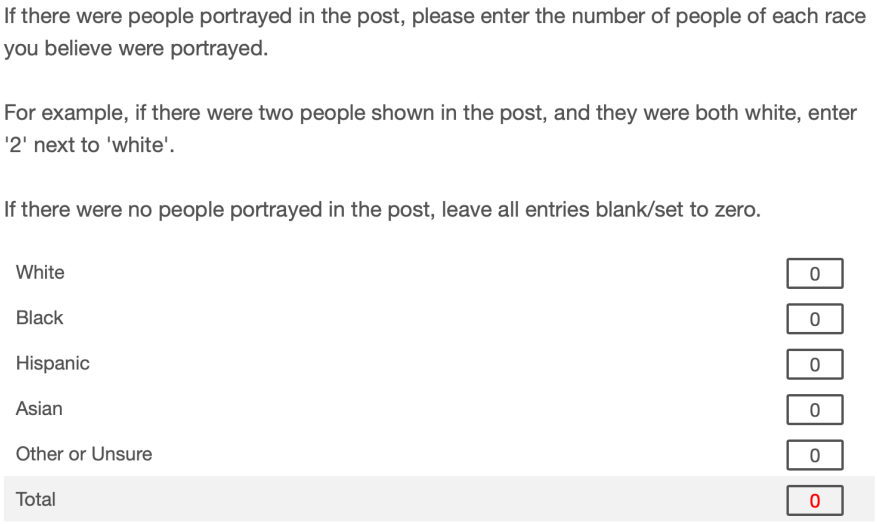

Fig. 19. Question on race of video subjects

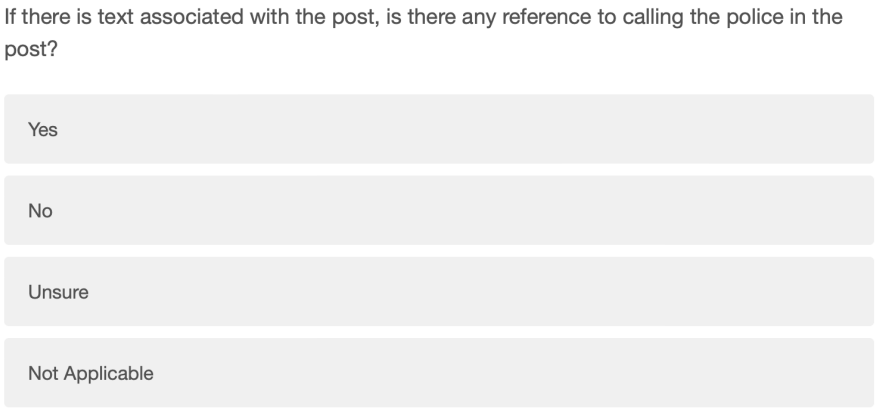

Fig. 20. Question on calling police.

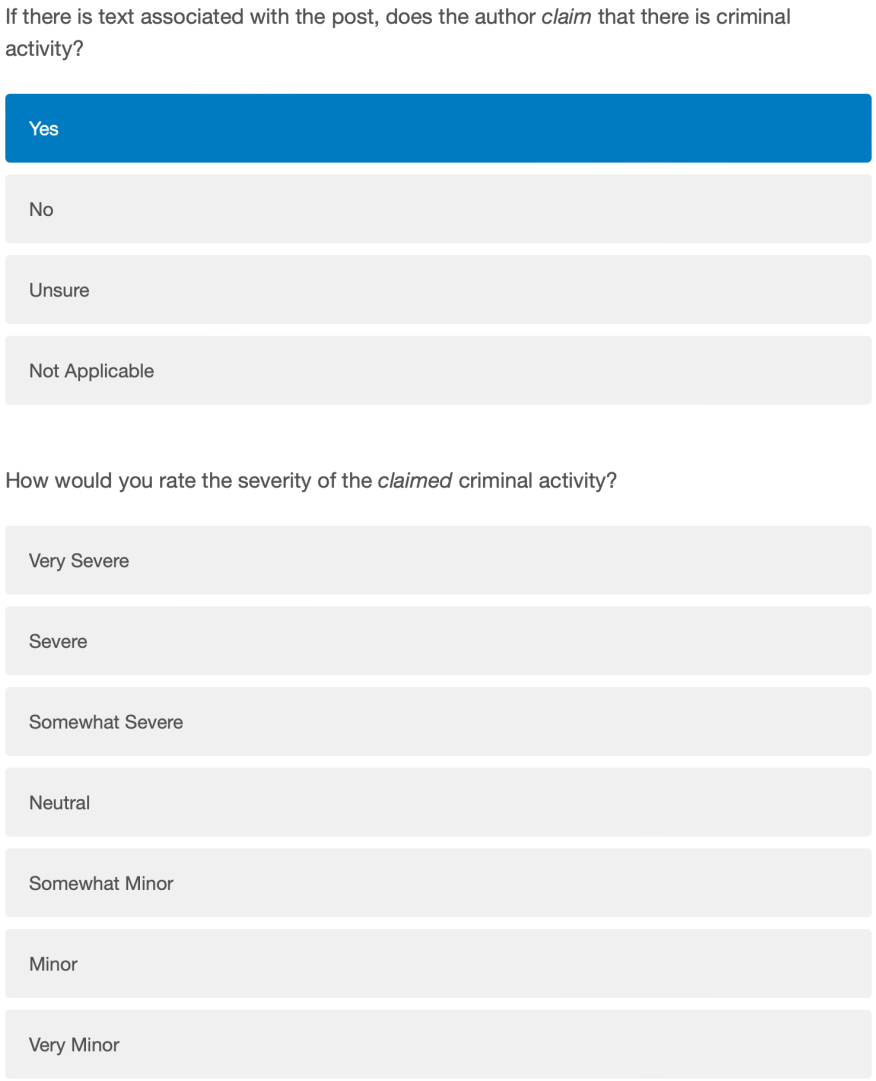

Fig. 21. Claimed criminal activity question

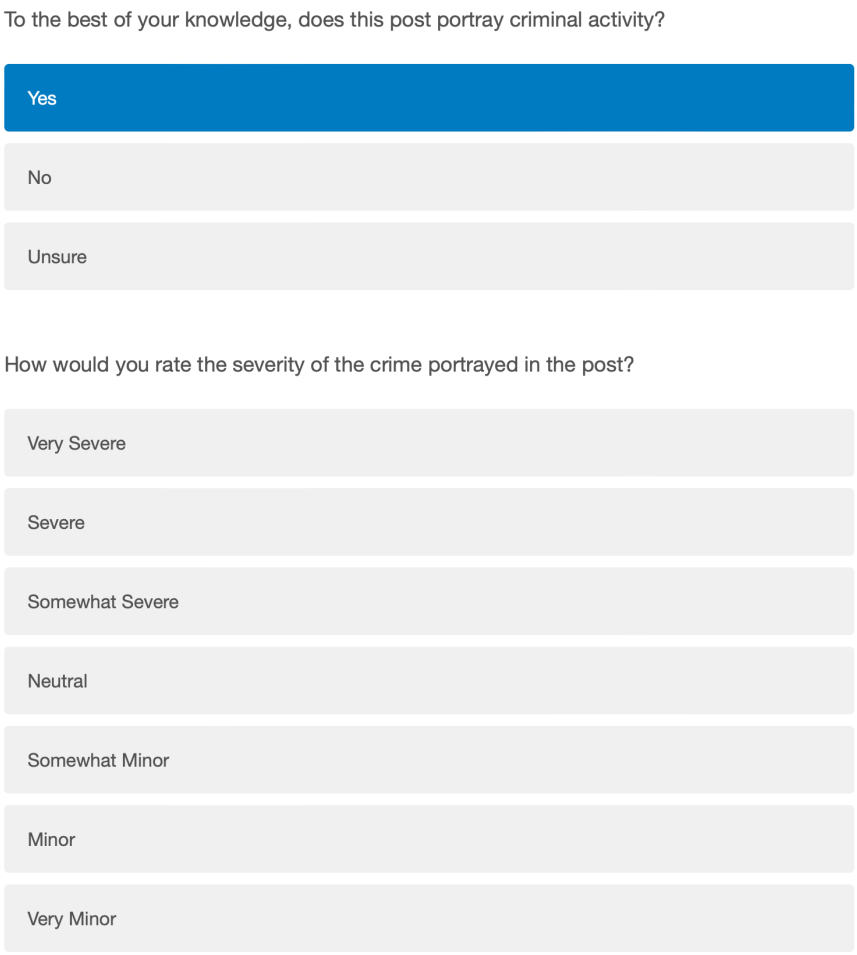

Fig. 22. Portrayed criminal activity question