Online Supporting Information

Social Networks and Protest Participation: Evidence from 130 Million Twitter Users

Jennifer M. Larson,[*] Jonathan Nagler,[†] Jonathan Ronen,[‡] and Joshua A. Tucker[§¶]

Contents

[*]Vanderbilt University Department of Political Science, 230 Appleton Place, Nashville, TN 37203. jennifer.larson@vanderbilt.edu

[†]NYU Department of Politics, 19 W. 4th St., New York, NY 10012. jonathan.nagler@nyu.edu

[‡]Berlin Institute for Medical Systems Biology, Max Delbrück Center for Molecular Medicine, Robert-Rössle Straße, Berlin, Germany 13125. yablee@gmail.com

[§]NYU Department of Politics, 19 W. 4th St., New York, NY 10012. joshua.tucker@nyu.edu

[¶]The data were collected by the NYU Social Media and Political Participation (SMaPP) laboratory (https://wp.nyu.edu/smapp/), of which Nagler and Tucker are co-Directors along with Richard Bonneau and John T. Jost. The SMaPP lab is supported by the INSPIRE program of the National Science Foundation (Award SES-1248077), the New York University Global Institute for Advanced Study, the Moore-Sloan Data Science Environment, and Dean Thomas Carew's Research Investment Fund at New York University.

1 Derivation of Hypotheses

The proofs rely on the contrast between sets of nodes that enter our dataset because they participated in the protest and nodes that enter our dataset because they were randomly sampled from among those who were eligible but did not participate. The next section explains why the comparison of these two differently-sampled sets is meaningful for our case and lays out the assumptions that establish exactly when these comparisons would be useful for future studies as well.

1.1 Setup and Working Assumptions

We begin by identifying a set of Twitter users of interest N from among the full universe of Twitter users $\mathcal{N}$. Here, N is defined to be set of Twitter users who were eligible to participate in the Charlie Hebdo protest. For our empirical analyses, we operationalize this definition with hashtag use and geolocation: N is the set of individuals who Tweeted with one of the hashtags pertaining to Charlie Hebdo during the time of the protest (for whom Charlie Hebdo was salient), and who were in Paris during the time of the protest (they were geographically near the protest site and had geolocation activated).

Our focus will be on a small slice of the Twitter universe's network: the network measured outward two degrees from individuals in N.

Definition 1 (2-degree network)**.** *For a set of nodes A contained within universe of nodes $\mathcal{A}$, the* ***2-degree network*** *of A is the egocentric network centered on egos A and measured out to geodesic distance 2 from any node in A.*

In other words, the 2-degree network of our eligible Twitter users N contains as nodes everyone in N, everyone to whom they are connected in $\mathcal{N}$, and everyone to whom their connections are connected in $\mathcal{N}$, and contains all links in paths of length 1 or 2 from anyone in N. Call the 2-degree network of N "G."

Our approach is to collect two subsets of N, each according to a different rule, and compare the subnetworks of G that are induced by the subsets *in order to learn something about the collection rules.* Here we lay out the necessary assumptions and logic underlying these comparisons.

Let P and C be two subsets of N. Our hypotheses pertain so long as subsets are small relative to N and are similar in size.[1]

[1]The reason for these assumptions will be clear below. As a brief preview, if $P \cup C = N$, then the

Assumption 1: *P and C are the same size. That is,*

$$|P| = |C|.$$

Assumption 2: *The size of P and the size of C are small relative to the size of N. That is, at a minimum,*

$$|P \cup C| < |N|.$$

The greater the difference between $|P \cup C|$ and $|N|$, the starker our results will be. Finally, our focus will be on differences in the networks among those in P and those in C, so some labels will be useful:

Definition 2. *Call g^P and g^C the 2-degree subnetworks of G induced by P and C, respectively.*

Each of g^P and g^C is a 2-degree network centered around egos P and C, respectively, and is contained in the broader network G.

Now let's consider different ways that elements of a subset of N can be chosen. To begin, let nodes posses exogenous, binary features called "attributes." For any attribute, a node either does or does not possess it.[2] Now we can define two collection rules by which subsets can be selected from the larger set:

Definition 3 (Random collection rule)**.** *Elements in subset $S \subset A$ are selected by a* ***random collection rule*** *if each element $s \in S$ was selected uniformly at random from A.*[3]

A subset collected via a random collection rule was formed by random selection; alternatively, a subset can be formed because nodes possess some attribute of interest:

---

fact that membership in P introduces dependencies between nodes in P with respect to the network will mechanically introduce dependencies between nodes in C with respect to the network. If N is large relative to $P \cup C$, then dependencies introduced by membership in P do not mechanically introduce dependencies in a randomly drawn C. In other words, these hypotheses could not be tested with samples of P and C such that $P \cup C = N$. The derivation of hypotheses is cleanest when we assume $|P| = |C|$, but the hypotheses are still valid tests when $|P| < |C|$ (so long as $|P \cup C|$ is still small relative to $|N|$.

[2]Clearly there can be interesting attributes of nodes that are not binary. Continuous attributes, such as length of time that a user's Twitter account has been open, can be converted to binary ones with a cutoff rule, such as "Twitter account older than 5 years." The results do not require that the attribute be binary, but their presentation is simplified with this assumption. Our attribute of interest, protest attendance, is binary.

[3]The key part of this definition is the randomness, which means that selection is orthogonal to any network feature. Whether elements are drawn from the whole set with probability $\frac{1}{|A|}$ or from a part that excludes some other subset S', resulting in probability $\frac{1}{|A \backslash S'|}$, is unimportant.

Definition 4 (Attribute collection rule)**.** *Elements in subset $S \subset A$ are selected by an **attribute collection rule** if each element $s \in S$ was selected from A because it possessed a certain attribute.*

Now let's assign a collection rule to each of the subsets defined above. Suppose subset $C \subset N$ is formed via the random collection rule, and $P \subset N$ is formed via an attribute collection rule.

Lastly, we need a few more definitions to set up our analyses and establish scope. Call $S^{a=1}$ the subset collected by including all nodes that possess attribute a, $S^{a=0}$ the complement subset that lacks attribute a, and S^r a subset collected with the random rule that is the same size as $S^{a=1}$. The fact that $|S^r| = |S^{a=1}|$ is important for the analyses that follow.[4]

Definition 5. *Call an attribute "rare" if substantially fewer than half of the members of the set of interest possess it. Specifically, $|S^{a=1}| << |S^{a=0}|$.*

Our analyses will focus on rare attributes.[5] This assumption allows us to collect both the full set of nodes that possess the attribute as well as a random set of the same size while still respecting the requirement that $|S^{a=1} \cup S^r| < |N|$.

Ultimately, we are interested in learning whether an attribute is meaningful to the network—whether it helps to delineates a meaningful social group or not. Let $N_i(g)$ refer to i's neighbors in network g and define the following possible characteristic of an attribute:

Definition 6. *Call an attribute "socially irrelevant" if the attribute is orthogonal to nodes' links. In other words, on average, for $i \in S^{a=1}$,*

$$\#N_i(g^{S^{a=1}}) = \#N_i(g^{S^r})$$

where S^r is a random subset of N the same size as $S^{a=1}$.

An attribute is socially irrelevant if an attribute-possessing node's neighbors are just as likely to be captured in a subset containing everyone possessing that attribute as they are to be captured in a randomly-chosen subset. That is, if collecting on the attribute makes it no more likely that a node with that attribute's neighbor is also collected than collecting at random does, then that attribute is socially irrelevant. For instance, we might imagine that

[4] Again, this assumption simplifies things but is not strictly necessary; the necessary assumption is that $|S^r| \geqslant |S^{a=1}|$ and $|S^r \cup S^{a=1}| < |N|$.

[5] Assuming rareness simplifies many of the comparisons below, and is fitting for our attribute of interest, protest attendance, which is rare among even eligible Twitter users in Paris.

the attribute "having the letter 't' appear at least three times in a Twitter user's username" is socially irrelevant. If we collected everyone that possesses this attribute in the set N, we probably have not collected a meaningful social group. Collecting with this attribute was as if we collected at random; the resulting social group would be equally without meaning. Our interest is in separating social irrelevance from social relevance:

Definition 7. *Call an attribute "socially relevant" if nodes are more likely than chance to be linked to other nodes sharing that attribute. That is, on average, for $i \in S^{a=1}$,*

$$\#N_i(g^{S^{a=1}}) > \#N_i(g^{S^r})$$

where S^r is a random subset of N the same size as $S^{a=1}$.

An attribute is socially relevant if a person is more likely to have neighbors who possess that attribute than chance. Collecting on a socially relevant attribute increases the likelihood that a node's neighbors are also included in the collection relative to collecting at random. An attribute can vary in its social relevance based how much more likely than chance a person is to have ties to others sharing that attribute. It will be useful to refer to the strongest possible form of social relevance:

Definition 8. *Call an attribute "perfectly socially relevant" if nodes are* ***only*** *linked to other nodes sharing that attribute. Specifically,*

$$\forall i \in S^{a=1}, \forall j \in N_i(g^{S^{a=1}}), j \in S^{a=1}.$$

When an attribute is perfectly socially relevant, all attribute-possessing nodes' neighbors also possess that attribute.

Our interest is in determining whether a particular rare attribute is socially relevant. Specifically, our manuscript presents a theory for why the process that leads up to a protest results in protest attendance being correlated across links—people exposed to others who intend to protest are more likely to protest themselves. Stated a different way, our manuscript presents a theory for why the attribute "attended the Charlie Hebdo protest" is socially relevant. We want to derive hypotheses that would be true of subsets of nodes collected on a rare, socially relevant attribute relative to a subset at least the same size collected at random.

The specific hypotheses are derived in the next sections, but the logic underlying them is as follows. For clarity, consider a rare attribute $a_{perfect}$ that is *perfectly* socially relevant.

Call the subset of nodes that possess this attribute $P_{perfect}$. By definition, individuals in $P_{perfect}$ only have social ties to others in $P_{perfect}$. Now also consider a subset of nodes C randomly selected from the rest of the set of interest, $N \backslash P_{perfect}$, which is the same size as the subset of those that possess the attribute, i.e. $|C| = |P_{perfect}|$.

The 2-degree subnetworks generated by these two subsets will differ in a few key ways. To see one, focus on the nodes in N to whom nodes in each of these subsets can be connected. Because $P_{perfect}$ was collected based on a perfectly socially relevant attribute, the set of nodes to whom nodes in this subset can be connected is also $P_{perfect}$. All ties of individuals in $P_{perfect}$ are also contained within $P_{perfect}$. Hence, the proportion of neighbors of any $i \in P_{perfect}$ that are also contained in $P_{perfect}$ is 1. Written formally,

$$\frac{1}{|P_{perfect}|} \sum_{i \in P_{perfect}} \frac{\#\{j | j \in P_{perfect}, j \in N_i(g^{P_{perfect}})\}}{\#\{j | j \in N_i(g^{P_{perfect}})\}} = 1.$$

The same is *not* true of the randomly selected subset C; individuals in C can be connected to any node in $N \backslash P_{perfect}$. Collecting only $|C|$ of them at random means that the probability of collecting any one neighbor of $i \in C$ in the subset C is

$$\frac{|C| - 1}{|N \backslash P_{perfect}| - 1},$$

which, by the assumption that $|C \cup P_{perfect}| < |N|$, is less than one.[6] Hence, the proportion of neighbors of anyone in C that are also contained in C is, in expectation, less than one, and by implication, less than the proportion of neighbors of anyone in $P_{perfect}$ also contained in $P_{perfect}$.

$$\frac{1}{|P_{perfect}|} \sum_{i \in P_{perfect}} \frac{\#\{j | j \in P_{perfect}, j \in N_i(g^{P_{perfect}})\}}{\#\{j | j \in N_i(g^{P_{perfect}})\}} > \frac{1}{|C|} \sum_{i \in C} \frac{\#\{j | j \in C, j \in N_i(g^C)\}}{\#\{j | j \in N_i(g^C)\}}.$$

(This is derived for the more general case of imperfectly socially relevant attributes below as Hypothesis 1.) Now contrast this scenario with a new one in which the subset collected

[6] Here is where the assumption that $|C \cup P_{perfect}| < |N|$ binds. If $C = N \backslash P_{perfect}$, then $\frac{|C|-1}{|N \backslash P_{perfect}|-1} = 1$ and none of the differences we report below would hold. The rarer the attribute (and hence the farther $|C \cup P_{perfect}|$ is from $|N|$, the starker the differences will be.

with an attribute collection rule uses a rare but socially *irrelevant* attribute $a_{irrelevant}$ (for instance, having at least three 't's in a username). Call this subset $P_{irrelevant}$, and once again consider a randomly drawn subset from the attribute-less nodes $N \backslash P_{irrelevant}$ of the same size, so that $|P_{irrelevant}| = |C|$. Because this attribute is socially irrelevant, the nodes to whom this subset may be linked are not confined to $P_{irrelevant}$; instead, anyone in $P_{irrelevant}$ could be linked to anyone else in N. This means the probability that the randomly selected set C contains any one neighbor of someone in C is

$$\frac{|C| - 1}{|N| - 1}, \tag{1}$$

and, since by definition a socially irrelevant attribute gives no information about social ties, the probability that the as-if-randomly selected set $P_{irrelevant}$ contains any one neighbor of someone in $P_{irrelevant}$ is

$$\frac{|P_{irrelevant}| - 1}{|N| - 1}. \tag{2}$$

Since, by assumption, $|C| = |P_{irrelevant}|$, ratio 1 is equal to ratio 2. That implies that the expected proportion of neighbors of those in $P_{irrelevant}$ that are contained in $P_{irrelevant}$ would not, on average, differ from the expected proportion of neighbors of those in C contained in C.[7] Hence, we would expect

$$\frac{1}{|P_{irrelevant}|} \sum_{i \in P_{irrelevant}} \frac{\#\{j | j \in P_{irrelevant}, j \in N_i(g^{P_{irrelevant}})\}}{\#\{j | j \in N_i(g^{P_{irrelevant}})\}} = \frac{1}{|C|} \sum_{i \in C} \frac{\#\{j | j \in C, j \in N_i(g^C)\}}{\#\{j | j \in N_i(g^C)\}}.$$

In short, if attribute a really is a perfectly socially relevant attribute, many differences in the subnetworks among $P_{perfect}$ and C, $g^{P_{perfect}}$ and g^C, will mechanically follow that would not if a were socially irrelevant. The same logic holds for imperfectly socially relevant attributes; the differences are in the same direction, and the greater the social relevance and the rarer the attribute, the starker the differences will be.

[7] And, if $|C| > |P|$, we would expect the proportion of neighborhoods of those in C that also are contained in C to be *higher* than that for $P_{irrelevant}$.

1.1.1 Assumptions about paths and valuations

We will assume that if two nodes are linked in the network (here if one follows the other), the link between them has non-zero weight:

ASSUMPTION 3: $\forall\ i, j \in N,$ *if* $d_{i,j} = 1,$ *then* $s_{i,j} > 0.$

This assumption implies that if $d_{i,j} < \infty$, then $s_{i,j} > 0$: if there is a path from i to j in network G, then the strength of that path is non-zero. This assumption is plausible given that the creation of ties in a social network requires some action, which probably would not be undertaken if the strength were exactly zero. This assumption simplifies the statement of results below, though is not strictly necessary. We could allow zero-strength ties and, so long as the distribution of tie strength values contains enough non-zero values and the zero values are distributed sufficiently widely throughout the network, the hypotheses would still follow.

Finally, the influence process in the article text implies that those who participated in a protest valued it highly enough to do so and those who did not participate did not value it highly enough. Hence, we will assume that the valuation held by those whom we observe to have protested is positive:

ASSUMPTION 4: $\forall\ i \in P, V_i > 0$ *and for all* $i \in C, V_i < 0.$

1.2 Intro to Proofs

The hypotheses we derive below and assess with our data test whether having participated in the Charlie Hebdo protest is a socially-relevant node-level attribute. Section 1.1 above defines social relevance and walks through the logic that underlies each of these hypotheses. In the subsections that follow, we present the reasoning for why that particular comparison follows from social relevance and the arguments of the theory presented in the manuscript. We begin by presenting and deriving a hypothesis that is additional to those in the manuscript text, and then derive the remaining four that do appear in the manuscript.

1.2.1 Density Hypothesis

An additional hypothesis follows from the model. This hypothesis is closely related to H1, so we only present H1 in the article text.

We should expect to observe greater "density" in the network among P than in the network among C. The subset of eligible participants that protested should have more of the total possible ties to others who protested than the same for a sample of eligible participants that did not protest. Specifically, we expect

$$\frac{1}{\#P}\sum_{i\in P}\frac{\#\{j|j\in N_i(g^P),j\in P\}}{\#P} > \frac{1}{\#C}\sum_{i\in C}\frac{\#\{j|j\in N_i(g^C),j\in C\}}{\#C}. \tag{3}$$

This expectation stipulates a difference in the distribution of individual densities: on average, those who protested should differ from a subset of eligible participants who did not protest with respect to the proportion of possible links present to others in the subset. Intuitively, if protesters influence their direct ties to protest, protesters should be highly connected to one another, more so than comparable people who apparently neither influenced nor were influenced by direct ties to attend. This can be stated as our first hypothesis:

DENSITY HYPOTHESIS: *The density of the protesters' network, g^P, is greater than the density of the eligible non-protesters' network, g^C (inequality 3).*

Of course, if protesters are the type of Twitter users who have more direct ties in general, we could observe the same relationship stipulated in the Density Hypothesis regardless of whether influence-by-exposure was at play. In this sense, the Density Hypothesis is a weak test and H1 is the stronger test.

1.2.2 Test of Density Hypothesis

	PROTESTERS	COMPARISON	t-STAT	SUPPORT
Density Hyp.	0.0044 (0.0075)	0.0003 (0.0009)	14.9 [log: 14.4]	Y
H1 Prop. ties within	0.0040 (0.0063)	0.0004 (0.0017)	15.0 [log: 11.0]	Y
H2 Prop. ties of ties within	0.0332 (0.0332)	0.0057 (0.0080)	22.2 [log: 35.0]	Y

Table 1: Replication of Table 1 from article, including density hypothesis test. Standard deviations in parentheses, t-statistic tests the null that values for protesters and the comparison set are the same. t-statistic on log-transformed data in square brackets. All three hypotheses are supported with high statistical confidence.

1.3 Proof of Density Hypothesis

Proof. We will show that the influence process described in the article text implies a difference in the 2-degree subnetwork induced by protesters compared to the 2-degree subnetwork induced by a random sample of non-protesters in terms of density and the extent of ties to others in the sample. Consider a random sample S of N that is small relative to N and the 2-degree subnetwork g^S induced by S. For an individual $i \in S$, there is an expected number of his ties, $N_i(G)$, that are to others sampled in S. The proportion of other individuals in S to whom i will have direct ties is $within_{i,S} = \#\{j|j \in S, j \in N_i(G)\}/\#S$, where $within$ is determined by G. Call the expected value of $within_{i,S}$ over $i \in S$

$$within_S = \frac{1}{\#S}\sum_{i \in S} within_{i,S}. \tag{4}$$

If P and C were random samples from N, then for any node in P or C, $within_P = within_C = within_S$ given G. Now consider the consequences of the influence process described in the article text. A node in G that is directly tied to another is influenced by him to join the protest if he values it highly enough. For any node $i \in P$, by virtue of attending, $V_i > 0$. Therefore, according to the influence process, if one node is labeled "protester," a node to which it is directly tied is more likely to be labeled "protester." Consequently, $within_P > within_C$. Since C and P are small relative to N, a random sample from N that discards members of P will be close to a random sample of P without replacement.[8] Furthermore, V_i is independent of V_j when $V_j < 0$. Then the result follows: $within_P > within_C$. This gives the Density Hypothesis: on average, $within_P > within_C$, e.g.

$$\frac{1}{\#P}\sum_{i \in P}\frac{\#\{j|j \in N_i(g^P), j \in P\}}{\#P} > \frac{1}{\#C}\sum_{i \in C}\frac{\#\{j|j \in N_i(g^C), j \in C\}}{\#C}. \tag{5}$$

□

1.4 Proof of Hypothesis 1

Proof. Hypothesis 1 relies on similar logic to the Density Hypothesis. Just as a random sample from a network should contain some proportion of each sampled node's direct ties on average, so should a random sample contain some proportion of each sampled node's neighborhood on average. Call that expected proportion $within^{neighb}_{i,S} = \#\{j|j \in S, j \in N_i(G)\}/\#\{j|j \in N_i(G)\}$. Because of the influence process, a node with label "protester" is more likely to have neighbors with label "protester." That is, the neighborhood of any node in P is more likely to contain other members of P than the neighborhood around a node in

[8] As the set of nodes to discard from a random sample becomes vanishingly small relative to the population from which a sample is drawn, the sample approaches a random sample without replacement.

S would be likely to contain another member of S. Call the expected value of $within_{i,S}^{neighb}$ over $i \in S$

$$within_S^{neighb} = \frac{1}{\#S} \sum_{i \in S} within_{i,S}^{neighb}. \tag{6}$$

Then we have $within_P^{neighb} > within_S^{neighb}$. By the same logic as above, since C is small, the random sample C from $N \backslash P$ approximates a random sample from N without replacement. Since V_i is independent of V_j when $V_j < 0$, the result follows: $within_P^{neighb} > within_C^{neighb}$, e.g.

$$\frac{1}{\#P} \sum_{i \in P} \frac{\#\{j | j \in P, j \in N_i(g^P)\}}{\#\{j | j \in N_i(g^P)\}} > \frac{1}{\#C} \sum_{i \in C} \frac{\#\{j | j \in C, j \in N_i(g^C)\}}{\#\{j | j \in N_i(g^C)\}}. \tag{7}$$

□

1.5 Proof of Hypothesis 2

Proof. The proof is almost identical to the proof of Hypothesis 2. Since $E_{i,j}(d_{i,j} = 2, s_{i,j} > 0) > 0$, and since the strength of all direct ties is non-zero, individuals are positively exposed to their ties of ties. Then by the same logic as above, on average, nodes in P have a larger proportion of their set of ties of ties included in the set P than nodes in C have included in the set C. To be precise, call the expected proportion of ties of ties in a set $within_{i,S}^2 = \#\{j | j \in S, j \in N_i^2(G)\} / \#\{j | j \in N_i^2(G)\}$. The expected value over all i in a randomly drawn sample S is

$$within_S^2 = \frac{1}{\#S} \sum_{i \in S} within_{i,S}^2. \tag{8}$$

Then we have $within_P^2 > within_S^2$. By the same logic as above, since C is small, the random sample from C from $N \backslash P$ approximates a random sample from N without replacement. Since V_i is independent of V_j when $V_j < 0$, result follows: $within_P^2 > within_C^2$, e.g.

$$\frac{1}{\#P} \sum_{i \in P} \frac{\#\{j | j \in P, j \in N_i^2(g^P)\}}{\#\{j | j \in N_i^2(g^P)\}} > \frac{1}{\#C} \sum_{i \in C} \frac{\#\{j | j \in C, j \in N_i^2(g^C)\}}{\#\{j | j \in N_i^2(g^C)\}}. \tag{9}$$

□

1.6 Proof of Hypotheses 3 and 4

Proof. Hypothesis 3 and Hypothesis 4 both rely on the same argument about strong ties that is the first step of the proof. Let all ties in G be characterized as either "strong" or "weak" such that if a tie from i to j is "strong" and a tie from i to k is "weak," then that $s_{i,j} > s_{i,k}$. A randomly selected set S from N would contain some proportion of the strong ties of S:

$$within_{i,S}^{strong} = \frac{\#\{j|j \in S, j \in N_i(G), s_{i,j} = strong\}}{\#\{j|j \in N_i(G), s_{i,j} = strong\}}. \tag{10}$$

In expectation, this value in S would be:

$$within_{S}^{strong} = \frac{1}{\#S}\sum_{i\in S} within_{i,S}^{strong}. \tag{11}$$

Recall from the assumptions of the influence process that $E_{i,j}(d_{i,j} = 1, s_{i,j} > 0) >$, and that $E_{i,j}$ is increasing in $s_{i,j}$ when $d_{i,j} = 1$. Furthermore, V_i is increasing in $E_{i,j}$ so long as $V_j > 0$. Consequently, since $V_i > 0\ \forall i \in P$, $within_P^{strong} > within_S^{strong}$. By the logic above, C approximates a random sample without replacement, and so $within_P^{strong} > within_C^{strong}$. Now we can state the hypotheses in terms of two operationalizations of tie strength. First, for $ij \in G$, assume that $s_{i,j} = strong$ if there exists a k such that $ik \in G$ and either $jk \in G$ or $kj \in G$. By this assumption, in general, for any node i, any triangle in a network provides two strong ties incident to i: with respect to i, for any $j, k \in N$ such that $ij \in G, ik \in G$, and either $jk \in G$ or $kj \in G$, $s_{i,j} = strong$ and $s_{i,k} = strong$. To simplify notation, let $ijk \in G$ mean $ij \in G$, $ik \in G$, and either $jk \in G$ or $kj \in G$. Now, any $i, j, k \in G$ such that $ijk \in G$ is a set of i's strong ties. If a set S were drawn at random from N, there would be some proportion of a node's sets of strong ties that would entail others in S,

$$within_{i,S}^{strong^1} = \frac{\#\{ijk \in g^S|j \in S \text{ or } k \in S\}}{\#\{ijk \in g^S\}}, \tag{12}$$

which, in expectation, would be

$$within_{S}^{strong^1} = \frac{1}{\#S}\sum_{i\in S} within_{i,S}^{strong^1}. \tag{13}$$

Since influence induces correlations in the V_is of those connected by strong ties, then it must be that $within_P^{strong^1} > within_S^{strong^1}$. By the same sampling logic as above, it follows that

$within_P^{strong^1} > within_C^{strong^1}$ and we have Hypothesis 3:

$$\frac{1}{\#P}\sum_{i\in P}\frac{\#\{ijk \in g^P | j \in P \text{ or } k \in P\}}{\#\{ijk \in g^P\}} > \frac{1}{\#C}\sum_{i\in C}\frac{\#\{ijk \in g^C | j \in C \text{ or } k \in C\}}{\#\{ijk \in g^C\}}. \tag{14}$$

Second, assume that a tie $ij \in G$ is strong if $ji \in G$; that is, a tie is strong if it is reciprocated, and it is weak if it is not reciprocated. On average, there is a number of reciprocated ties that any node in N has in network G. In a random sample S, some number of $i \in S$'s reciprocated ties would be to other $j \in S$. Call this number $within_{i,S}^{strong^2} = \#\{j | ij \in G, ji \in G\}$, with expected value over S

$$\frac{1}{\#S}\sum_{i\in S} within_{i,S}^{strong^2}. \tag{15}$$

Now, *so long as* $\#P = \#S$, since influence induces correlations in the V_i of those connected by strong ties, $within_P^{strong^2} > within_S^{strong^2}$. And, by the same argument about sampling used the above derivations, so long as $\#P = \#C$, $within_P^{strong^2} > within_C^{strong^2}$. This gives us Hypothesis 4:

$$\frac{1}{\#P}\sum_{i\in P}\#\{j | ij \in g^P, ji \in g^P, j \in P\} > \frac{1}{\#C}\sum_{i\in C}\#\{j | ij \in g^C, ji \in g^C, j \in C\}. \tag{16}$$

The larger the difference between $\frac{1}{\#P}\sum_{i\in P}\#\{j | ij \in g^P, ji \in g^P, j \in P\}$ and $\frac{1}{\#C}\sum_{i\in C}\#\{j | ij \in g^C, ji \in g^C, j \in C\}$, the more likely the difference will appear even when $\#C > \#P$. In other words, a conservative test of this hypothesis will be the case in which $\#C > \#P$.

□

2 Additional Analyses and Robustness Checks

In the following sections, we consider alternative explanations that could undermine our conclusions.

2.1 Sensitivity of Descriptive Statistics

Protesters have more ties on average on Twitter than users in the Paris comparison set. They also have more than users in an additional comparison set drawn from France (described in more detail in the next section). Figure 1 shows the distribution of the number of friends for all three sets of Twitter users, and the same when the twenty largest values are excluded. Although the means compress slightly, protesters still have substantially more ties than either of the comparsion groups.

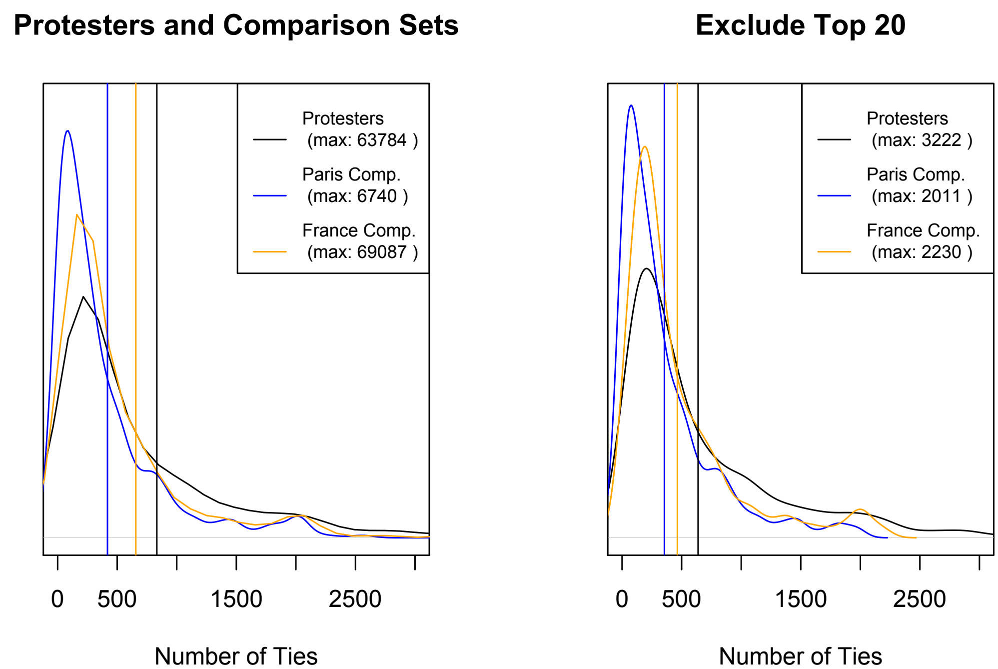

Figure 1: Distribution of the number of ties per user in the set of protesters, the Paris comparison set and the France comparison set. The left plot shows the raw density for all three; the right shows the densities for each, excluding the twenty largest values. Vertical lines plot the distributions' means.

Table 2 reports different summary statistics of the number of ties for users among protesters and the two comparison sets. The maximum number of ties in the protester set is substantially larger than the maximum in the Paris set, but not in the France set. However, these large values are more of an anomaly to the France comparison set than to the set of protesters: only 5% of the France comparison set have over $2,000$ ties, compared to 8% of the protesters. The median value of ties, and the mean excluding the top twenty values, are greater for protesters.

Figure 2 shows the same distribution for the number of ties of ties for all three groups. Once again, excluding the twenty largest values compresses the means slightly (displayed as vertical lines), but protesters have substantially more ties of ties on average (though the maximum value in the France comparison set is larger).

	Mean	Median	Max	Mean - top 20	Prop > 2k
Protesters	833	410	63784	638	0.082
Paris	418	226	6740	355	0.027
France	656	298	69087	464	0.051

Table 2: Values for the set of protesters, the Paris comparison set and the France comparison set. Mean - top 20 is the maximum value reported, excluding the 20 largest. Prop > 2k is the proportion of the sample with at least 2,000 ties.

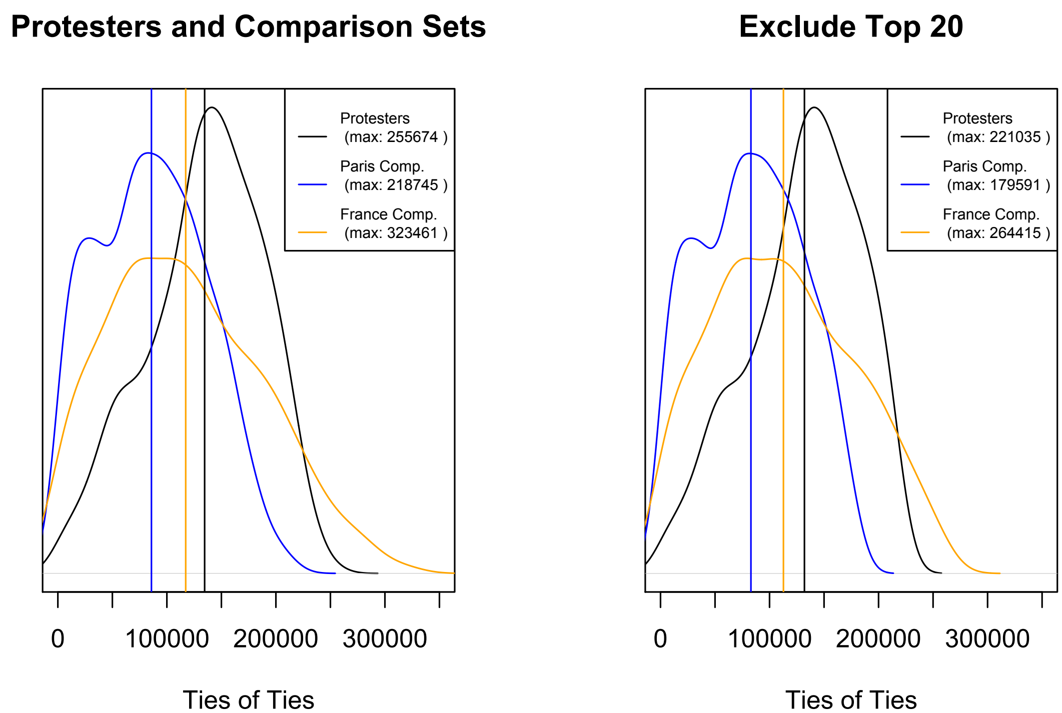

Figure 2: Distribution of the number of ties of ties per user in the set of protesters, the Paris comparison set and the France comparison set. The left plot shows the raw density for all three; the right shows the densities for each, excluding the twenty largest values. Vertical lines plot the distributions' means.

Table 3 repeats the comparisons of Table 2 for ties of ties. Protesters in general have more ties of ties, both on average, and at the median. The comparison holds when the twenty largest values are excluded. As with the number of ties, the France comparison set contains a higher maximum value of ties of ties, and here the proportion in the right tail is larger for the France comparison set. In other words, the extreme values of ties of ties are not greater among protesters, but the typical values– the mean and median– are.

	Mean	Median	Max	No Top 20	Prop > 2k
Protesters	134623	139944	255674	132007	0.102
Paris	85830	85808	218745	82936	0.008
France	117260	110651	323461	112747	0.139

Table 3: Comparisons of the set of protesters, the Paris comparison set and the France comparison set in terms of ties of ties. Mean - top 20 is the maximum value reported, excluding the 20 largest. Prop > 200k is the proportion of the sample with at least 200,000 ties of ties.

Figure 3 shows the distribution of transitivity for protesters compared to the same for the Paris and France comparison sets. Transitivity is computed as the ratio of weakly closed triangles to which a node is incident to the total number of possible triangles to which a node could be incident.

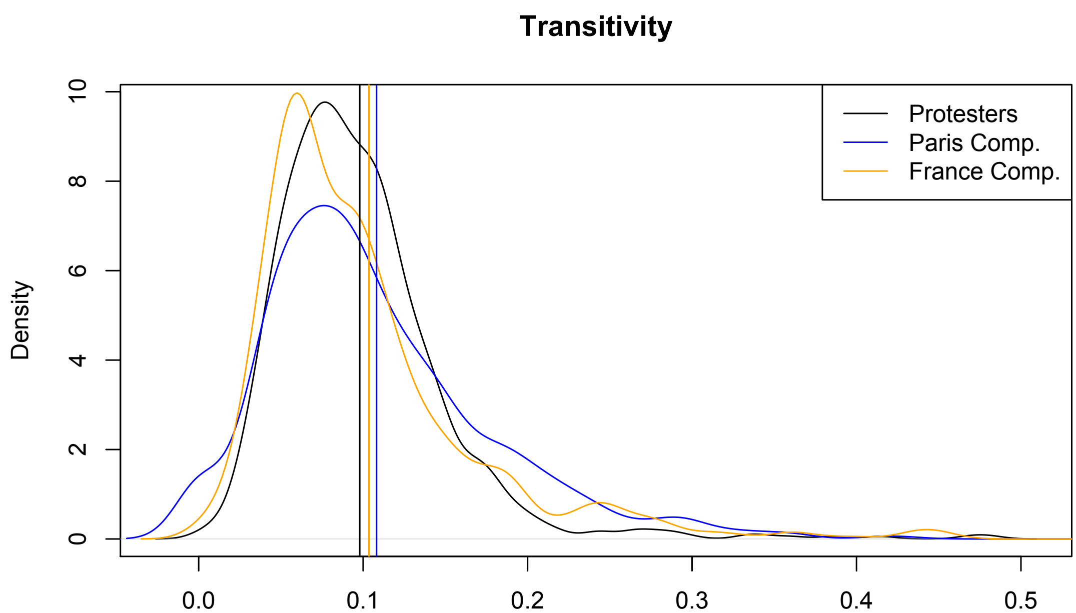

Figure 3: Distribution of transitivity for protesters, Paris comparison set and France comparison set. Vertical lines are means.

2.2 Robustness to France Comparison Group

Table 4 replicates the main results using a comparison group collected from Twitter users geo-located to be in France during the Charlie Hebdo protest but not present at the protest in Paris. These users were selected based on their location but *not* based on their hashtags. Unlike the other two samples, these users were not collected based on using the seven Charlie Hebdo hashtags (see manuscript section 3). This section shows that all results hold when using the comparison set from France instead of the comparison set from Paris.

	Protesters	FRANCE Comp.	t-stat	Reject
Prop. ties within	0.004 (0.008)	0.001 (0.002)	13.7 (log: 8.0)	Y
Prop. ties of ties within	0.033 (0.033)	0.004 (0.011)	22.9 (log: 40.5)	Y
Triangles within	0.011 (0.015)	0.002 (0.006)	15.0 (log: 8.1)	Y
Recip. within	3.1 (5.77)	.28 (.93)	13.2 (log: 12.6)	Y

Table 4: Hypothesis tests using the France comparison group. t-stat tests the null that the values are the same for the protesters and the comparison set. All hypotheses are supported.

Figure 4 shows the distribution of the proportion of ties within each of the groups for the protesters and both comparison sets. The right plot zooms in to the mass of the distribution. Protesters are a much more cohesive group: a much larger proportion of each protester's ties are to other protesters compared to the proportion of each member of the comparison set's ties that are to other members of each comparison set.

Protesters have many more ties within the group than either of the two comparison groups do, as Figure 7 shows.

The same comparisons hold for the proportion of ties of ties within each set, regardless of the comparison set. Figure 6 shows the distribution of the proportion of ties of ties within each group both in full (left) and zoomed in (right).

Finally, the proportion of triangles in the protest set containing other protesters is larger than the proportion of triangles in either of the comparison sets containing other individuals in the comparison sets. Figure 8 shows the distributions.

Figure 9 shows that while many users in all sets have no strong ties, a much higher proportion of users in the comparison sets have no strong ties..

It is reassuring that all results hold when we compare protesters to Twitter users in France. We are also reassured by the extent to which the Paris and the France comparison sets are similar to each other and are often more similar to each other than to the protesters.

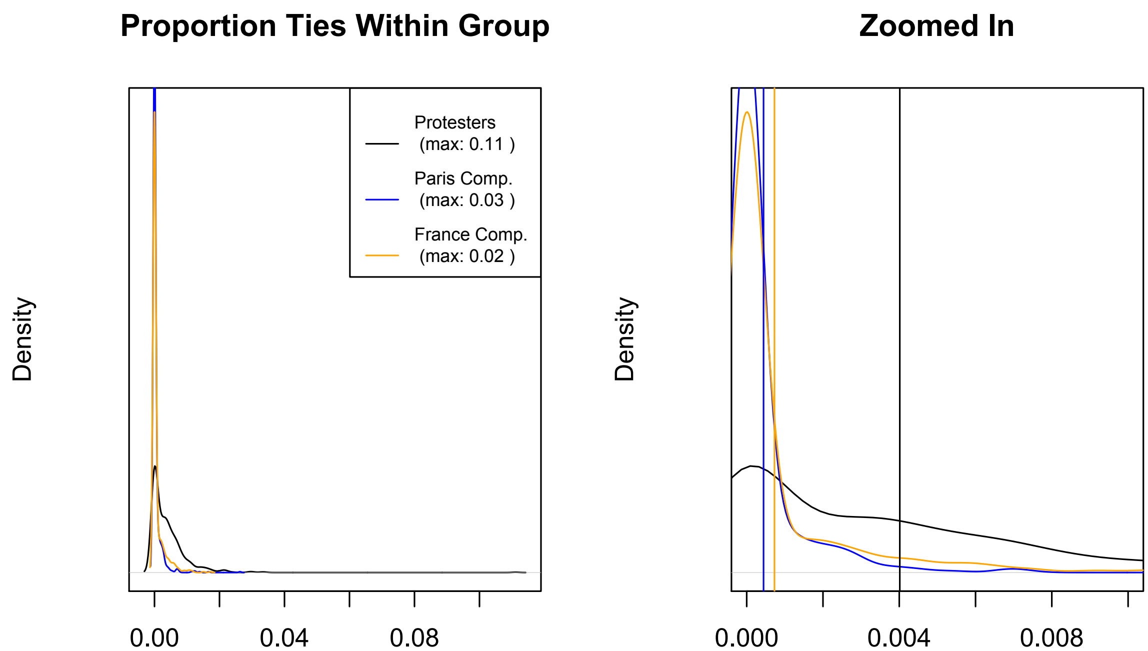

Figure 4: Proportion of each user's ties that are to other members of the relevant set–protesters to protesters, Paris comparison set to Paris comparison set, France comparison set to France comparison set.

2.3 Robustness to Omitting Verified Accounts

This section addresses the potential concern that our analyses include all ties on Twitter, including those to "verified accounts" which we would not expect to be important to an influence-by-exposure process. Here we pull all verified accounts and replicate our results. Table 5 replicates the main results, omitting all verified Twitter accounts. All comparisons are robust to excluding verified accounts.

	Protesters	Comparison	t-stat	Reject
Prop. ties within	0.004 (0.007)	0.001 (0.002)	14.0 (log: 10.6)	Y
Prop. ties of ties within	0.027 (0.028)	0.005 (0.008)	20.5 (log: 32.4)	Y
Triangles within	0.013 (0.020)	0.002 (0.007)	14.3 (log: 7.4)	Y
Recip. within.	2.8 (5.38)	0.15 (0.57)	13.5 (log: 13.8)	Y

Table 5: Hypothesis tests with verified accounts omitted. t-stat tests the null that the values are the same for the protesters and the comparison set.

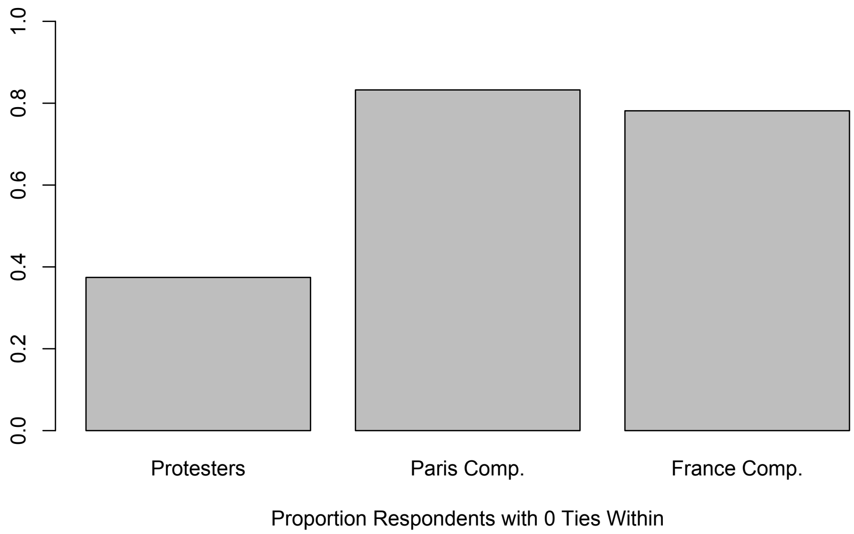

Figure 5: Proportion of respondents in each set– Protesters, Paris comparison, France comparison– that have no ties to any other user in that set.

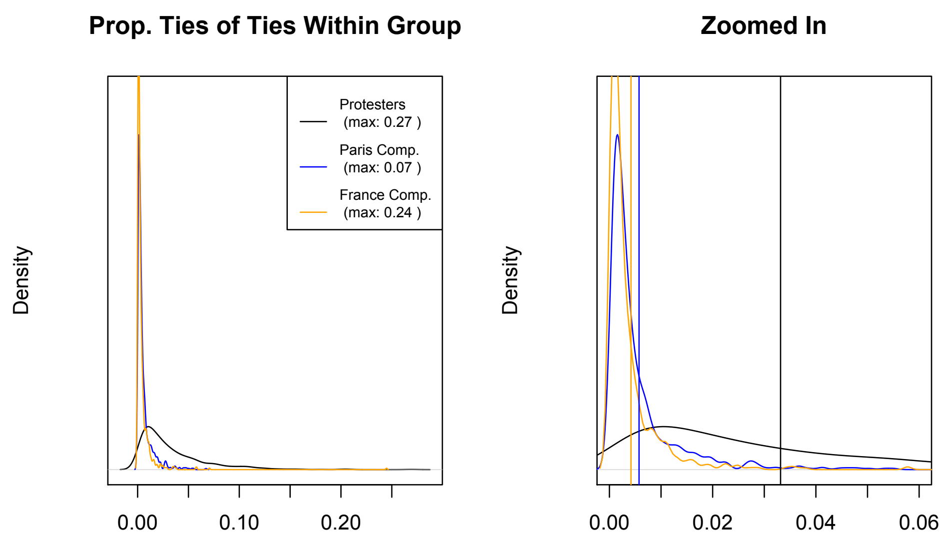

Figure 6: Proportion of each user's ties of ties that are to other members of the relevant set– protesters to protesters, Paris comparison set to Paris comparison set, France comparison set to France comparison set.

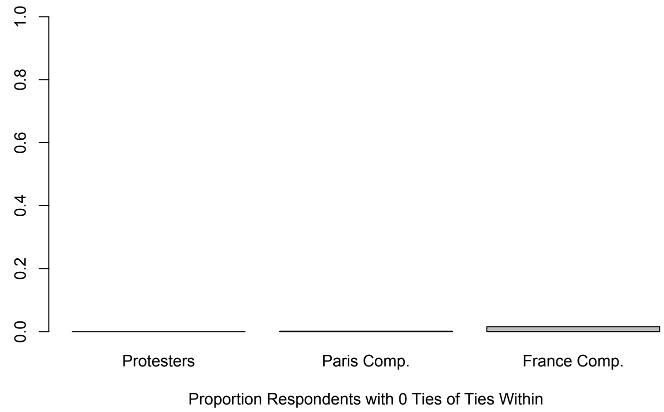

Figure 7: Proportion of respondents in each set– Protesters, Paris comparison, France comparison– that have no ties of ties to any other user in that set. The value is 0 for Protesters and the Paris comparison set, and 1.6% for the France comparison set.

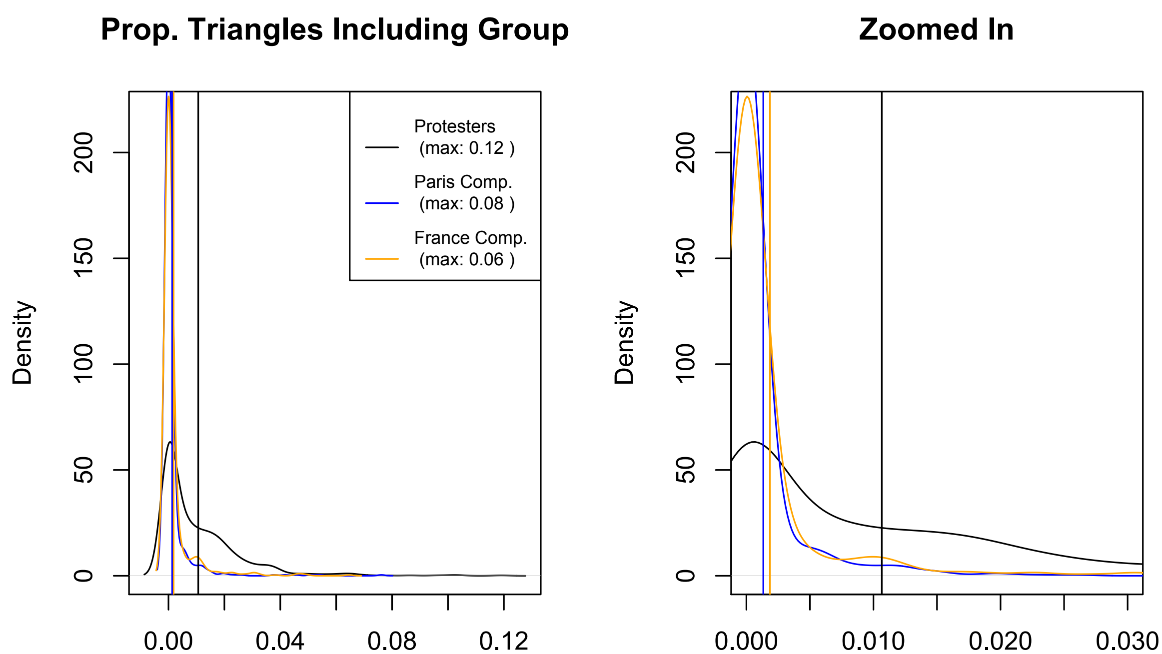

Figure 8: Proportion of respondents' triangles in each set– Protesters, Paris comparison, France comparison– that entail at least one other user in that set.

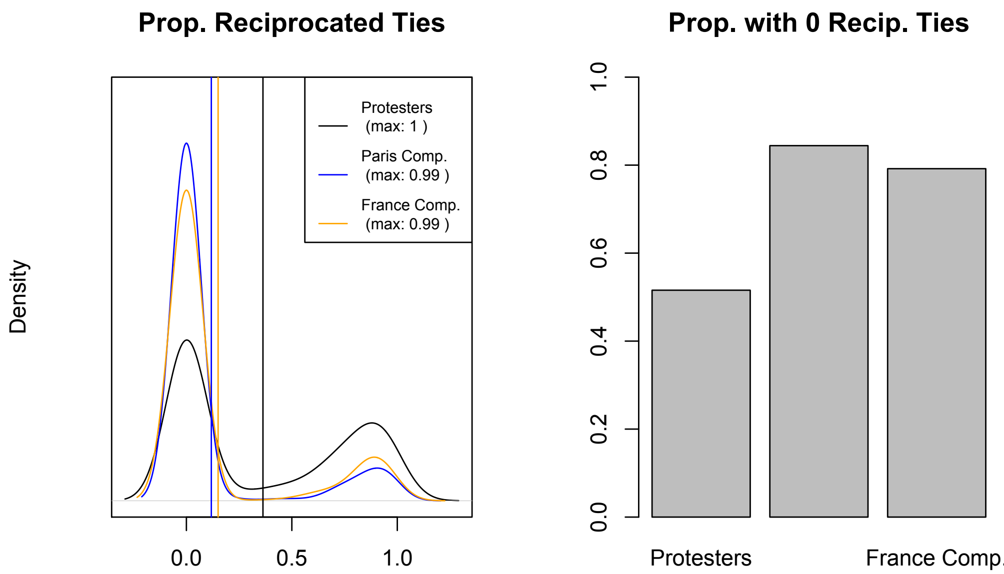

Figure 9: Proportion of respondents' strong ties in each set– Protesters, Paris comparison, France comparison.

Figures 10, 11, 12 and 13 replicate the results of the main article with verified accounts removed.

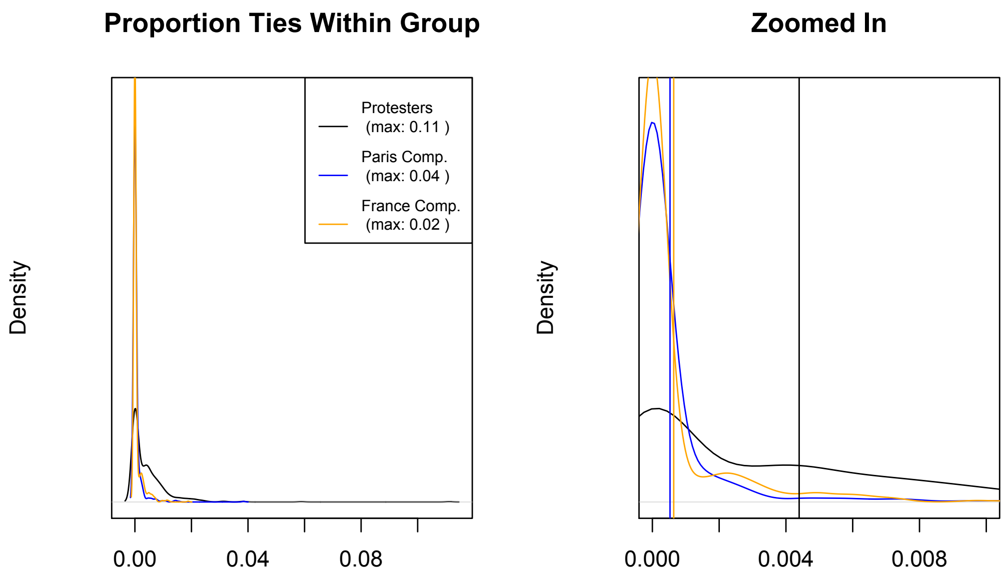

Figure 10: Proportion of ties among protesters and comparison sets with verified accounts omitted.

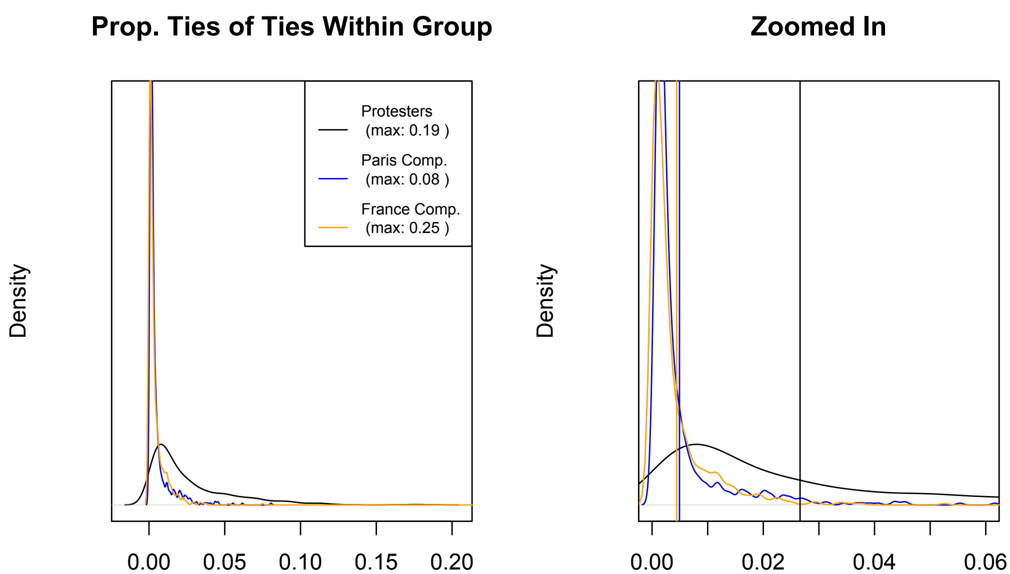

Figure 11: Proportion of ties of ties among protesters and among those in the comparison sets with verified accounts omitted.

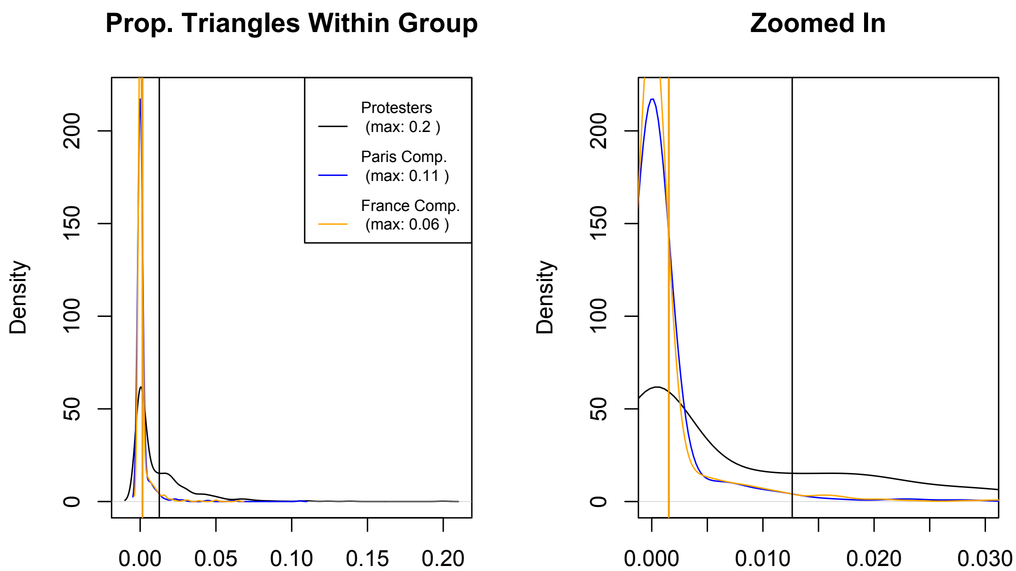

Figure 12: Proportion of triangles among protesters and among those in the comparison sets that entail others in the group, with verified accounts omitted.

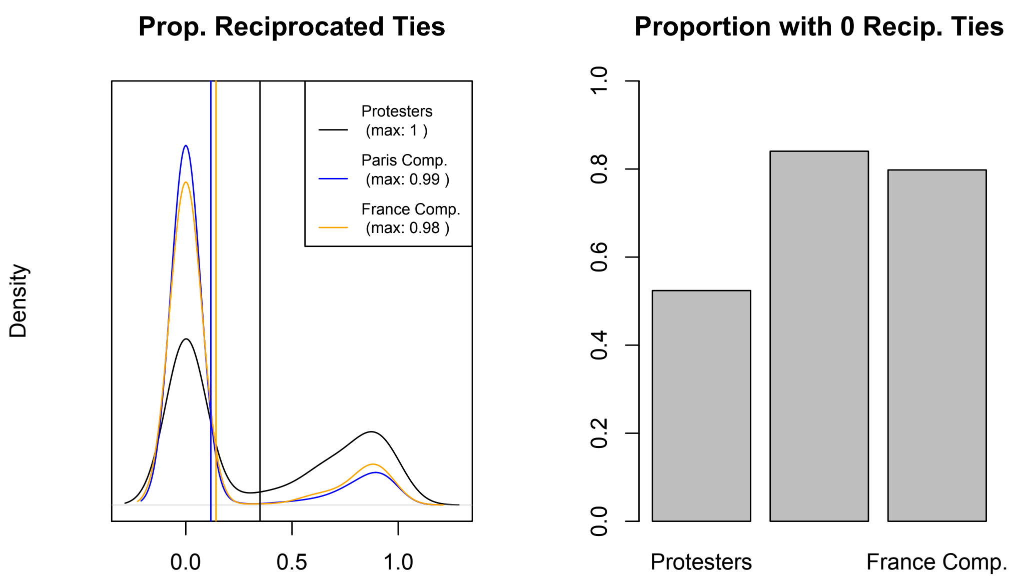

Figure 13: Proportion of strong ties, and incidence of strong ties among protesters and members of the comparison sets, with verified accounts omitted.

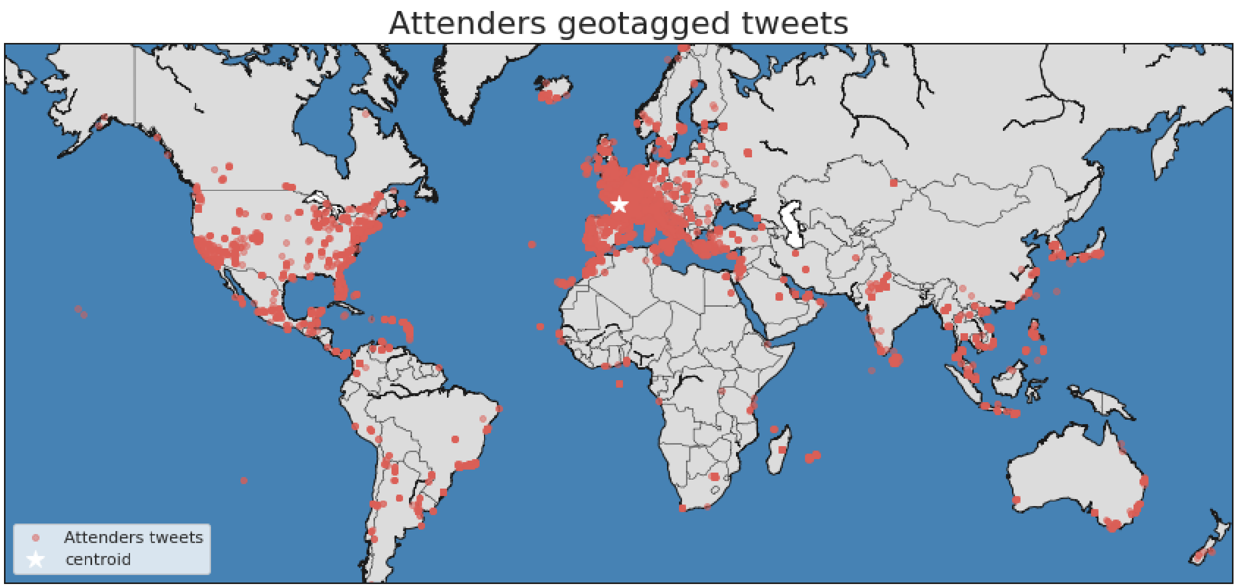

Figure 14: Geotagged tweets sent by protesters in the lifetime of their Twitter accounts

2.4 Mobility and Geolocation

This section presents some general information that speaks to the mobility of users in our samples and to their tendency to geolocate. Figure 14 displays the geographic distribution of tweets sent by protesters in the lifetime of their Twitter accounts. Figures 15 and 16 show the same for users in the Paris and France comparison sets, respectively.

Clearly, the protesters (labeled attenders) are not a more geographically concentrated group in terms of mobility; many Tweets are sent far from Paris. Furthermore, figure 17 constructs a measure of any geotagged Tweet's distance to the centroid of all geotagged Tweets. By this measure, protesters are a bit more mobile than members of either comparison set.

Finally, figure 18 shows the box plots of geotagging activity by user in each of the three sets. All three exhibit similar geotagging propensities. On average, users who attended the protest geocode 11.6% of their Tweets (sd=.14) and users in the Paris control geocode 10% of their Tweets (sd=.13).

2.5 Clustering by Accounts Followed

Figure 19 shows the results of a Latent Dirichlet Allocation decomposition based on accounts followed. If users who are more politically active follow certain political or news accounts to a

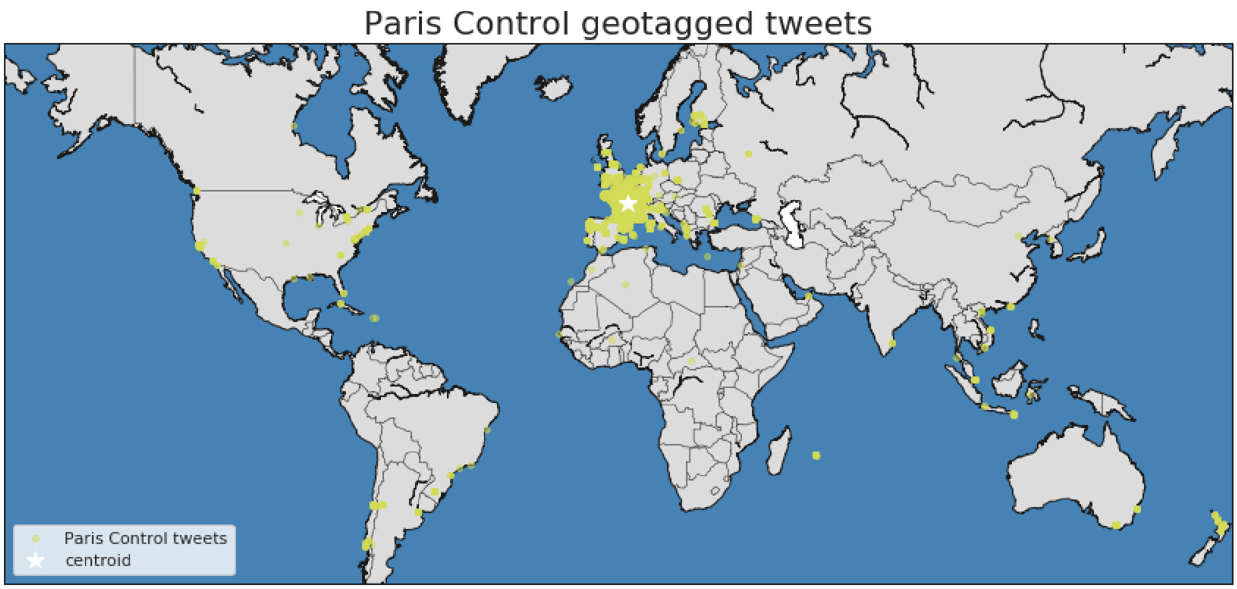

Figure 15: Geotagged tweets sent by users in the Paris comparison set in the lifetime of their Twitter accounts

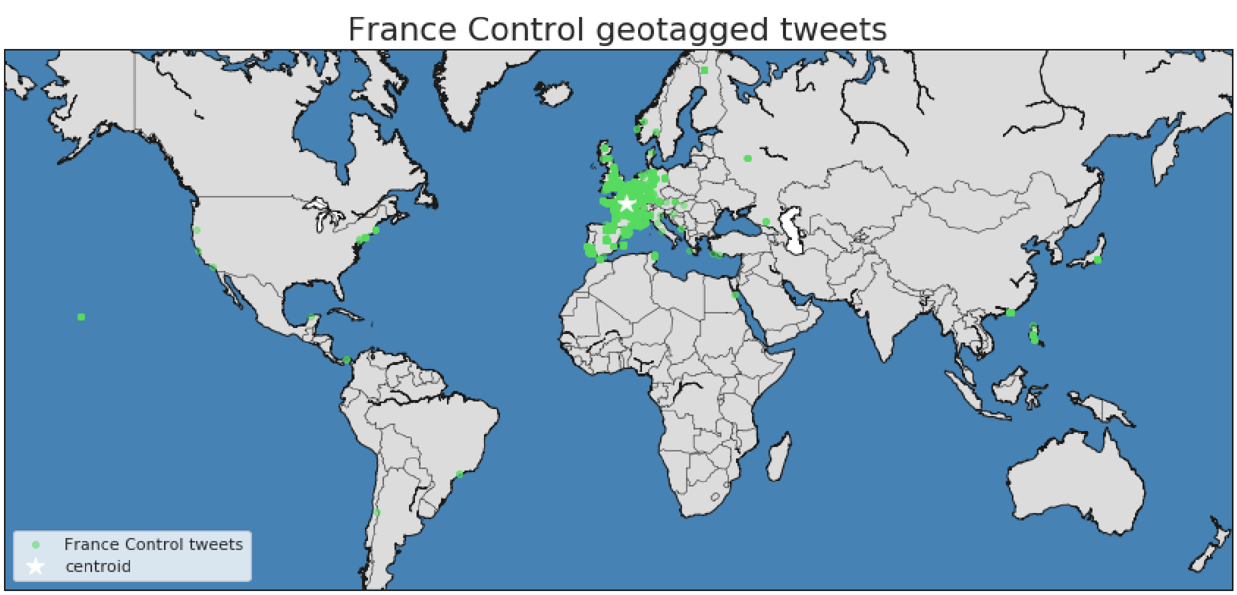

Figure 16: Geotagged tweets sent by users in the France comparison set in the lifetime of their Twitter accounts

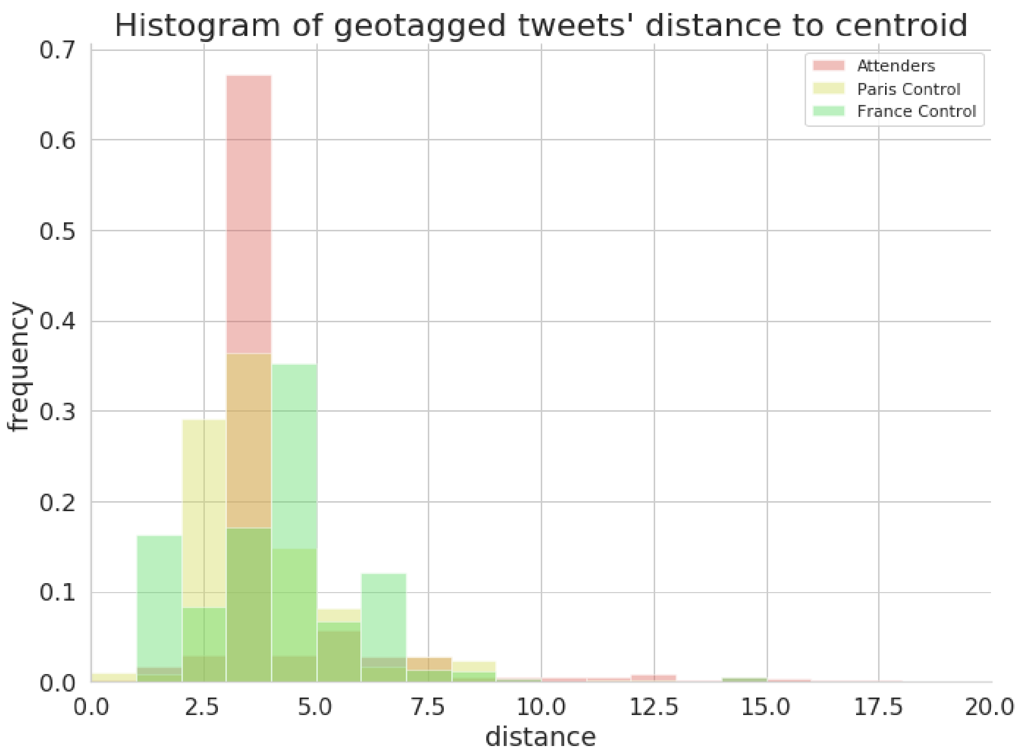

Figure 17: Distribution of the distance from geotagged Tweets sent in users' account lifetime from the centroid of these Tweets.

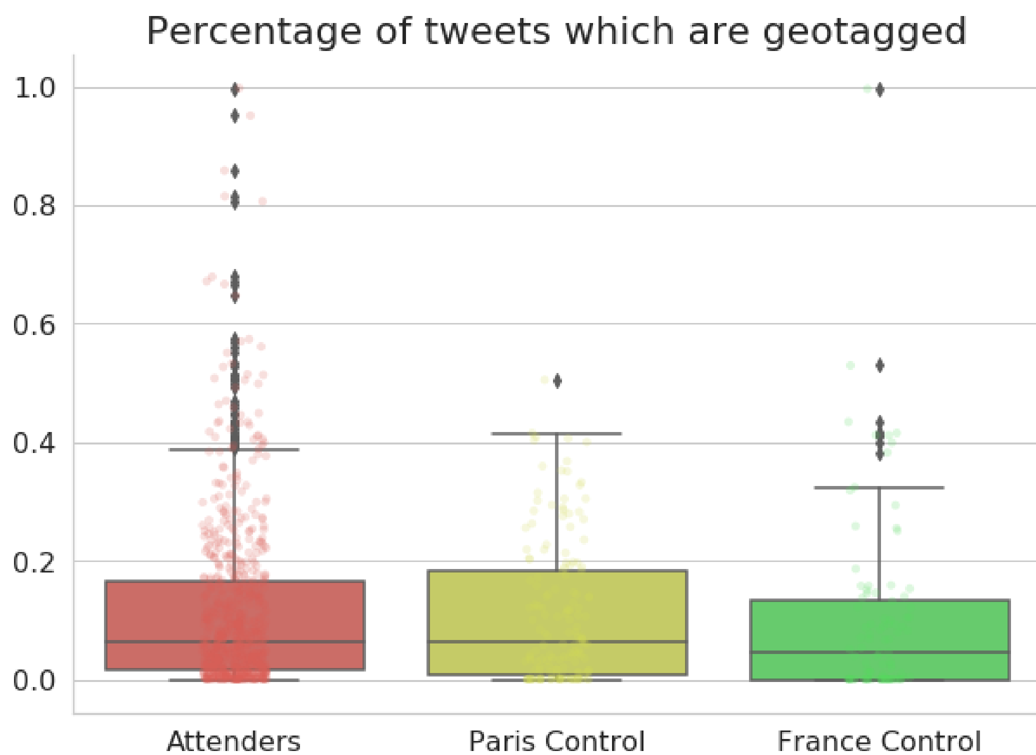

fig

Figure 18: Proportion of users' Tweets that are geotagged in the lifetime of that user's Twitter account

greater extent than users who are not politically active, and if protesters and non-protesters sort based on level of political activity, then we should observe clusters of protesters separate from clusters of non-protesters. Instead, Figure 19 shows no such clusters. We take this as suggestive evidence that there is no clear difference in level of political activity between protesters and the comparison set. If selection based on political activity is a threat to our inferences, we are likely controlling for political activity by using a comparison set that is similar to the protesters in level of political activity implied by their Twitter relationships.

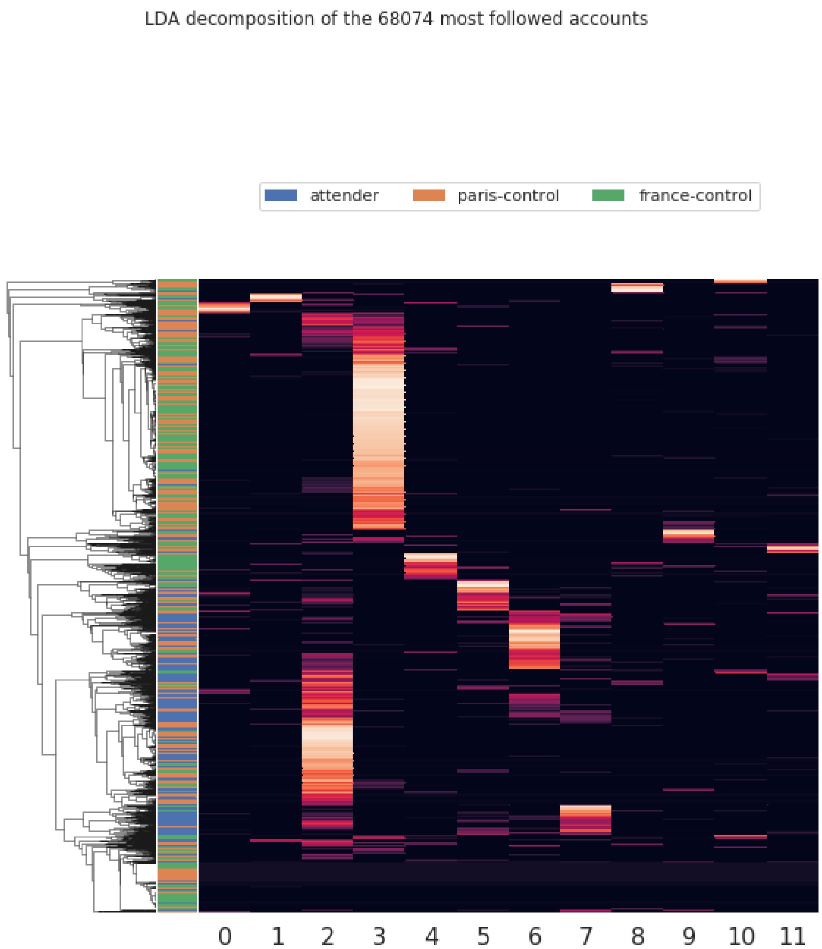

Figure 19: LDA decomposition of the 68,074 most followed accounts. Users do not cleanly separate into protester and non-protester clusters based on accounts that they follow.

2.6 Potential Baseline Differences

Here we consider potential differences between the Protesters and the Comparison set that may account for the differences in networks. First, we document the propensity to follow verified accounts. We see similar tendencies to follow verified accounts across the three groups, shown in Figure 20. To the extent that following verified accounts captures political interest or social consciousness (these accounts include news organizations and politicians), protesters are similar to the control sets.

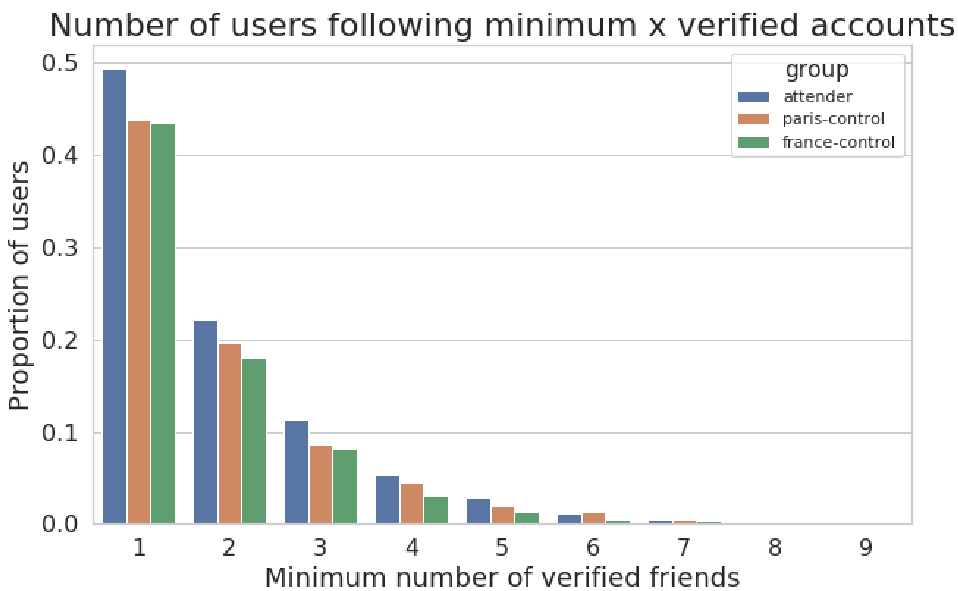

Figure 20: Number of users following at least x verified accounts.

Next, we look at the extent to which users may vary in their activity on Twitter. It could be that users with a certain gregarious or outgoing personality are more likely to be friends (and hence to follow each other on Twitter), so that the clustering we observe among Protesters is an artifact of this personality trait. To test for this, we look at all three sets' Twitter activity between the massacre and the protest, and in the days immediately after.

Figure 21 shows that the Protesters and the Comparison set have similar patterns and levels of Twitter activity, suggesting that Protesters are not merely more active on Twitter, or more outgoing (to the extent that this personality trait is captured by tweeting). Figure 22 shows Twitter activity specifically about Charlie Hebdo. These tweets use one of the hashtags about Charlie Hebdo. Here too, the Protesters and the Comparison set behave similarly, further casting doubt on a difference in level of interest in political topics or current events.

Note also that the protesters and users in the Paris comparison set are much more similar to one another than to the France comparison set. Selecting on the seven Charlie Hebdo hashtags resulted in a comparison set very similar to the protesters in terms of political interest, engagement with current events, sociability on Twitter, and gregariousness. Users in the France comparison set, which was not collected based on the hashtags, look very different in terms of these attributes.

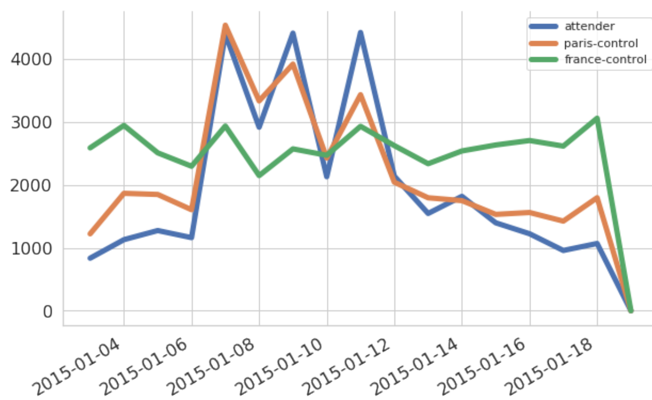

Figure 21: Twitter activity of three sets in window of time including protest.

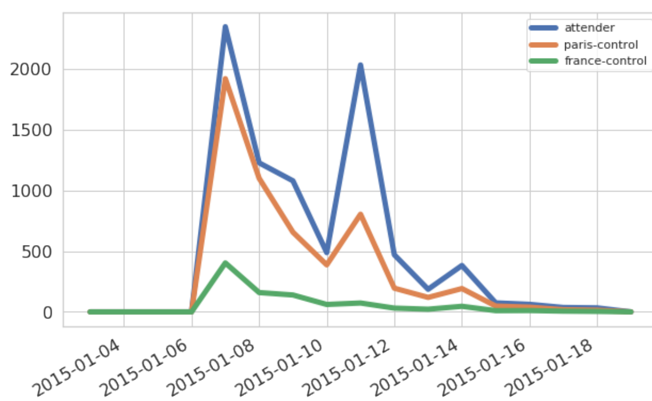

Figure 22: Twitter activity about Charlie Hebdo of three sets in window of time including protest.

Finally, we consider possible evidence of differences based on which users tweeted which of the seven Charlie Hebdo hashtags. Figure 23 shows that the protesters and the comparison set used the hashtags in very similar proportions. If different hashtags pick up different aspects of the issue of Charlie Hebdo, the distribution of interest in these aspects is very similar across our two sets.

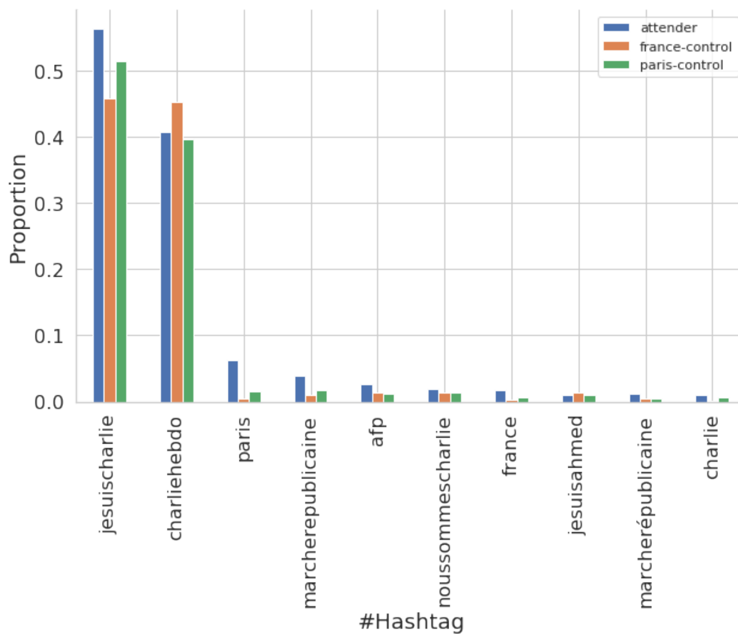

Figure 23: Distribution of hashtag use.

2.7 Robustness to following Protesters in both sets

Next we compare the extent to which protesters have ties to protesters to the extent to which non-protesters have ties to protesters. As expected, protesters have both a larger proportion of their ties and of their ties of ties to other protesters.

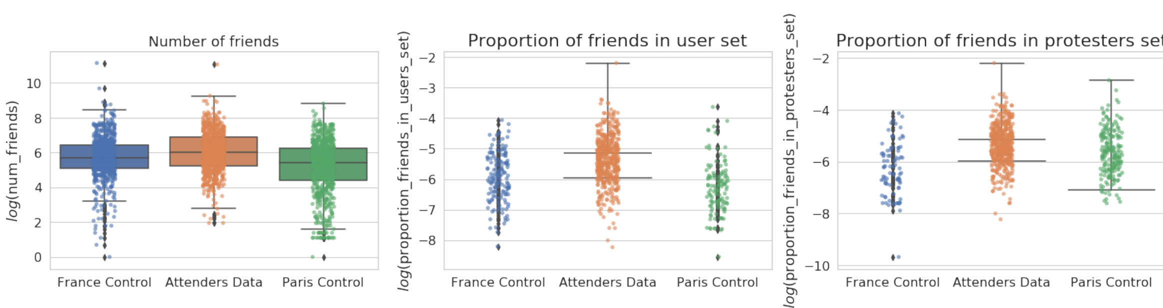

Figure 24: Proportion of ties that are protesters.

2.8 Did the Protest Cause the Network Ties?

Our analyses assume that the Twitter network is a measure of the underlying social network of potential protesters in the days leading up to the protest. Here we consider the possibility

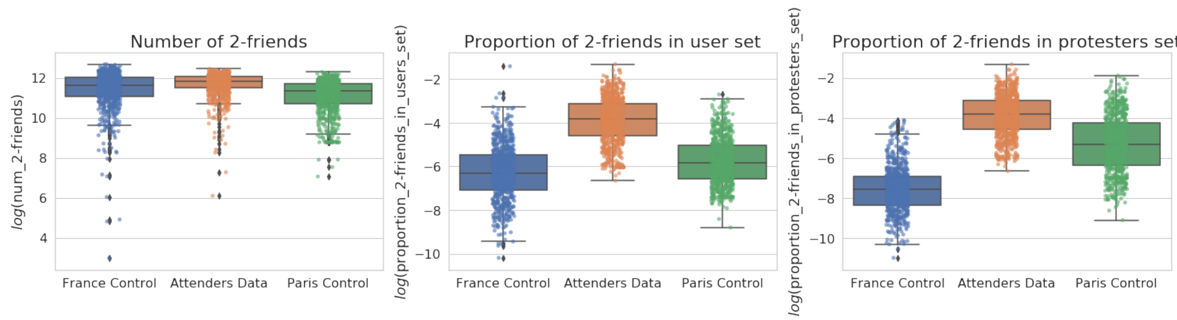

Figure 25: Proportion of ties of ties that are protesters.

that links between protesters were instead added to the network *after* the protest as a cause of protesting together. This would generate the results we observe but would lead us to misattribute the reason for them. While friendship ties (following relationships) are not timestamped, the order in which they were added is available from the Twitter API. We compare the rank of ties that are between protesters to the rank of ties that protesters have to other users. Figure 26 shows that the distribution of ranks for these two types of ties is very similar, making the creation of new protester-protester ties at one time following the protest unlikely.

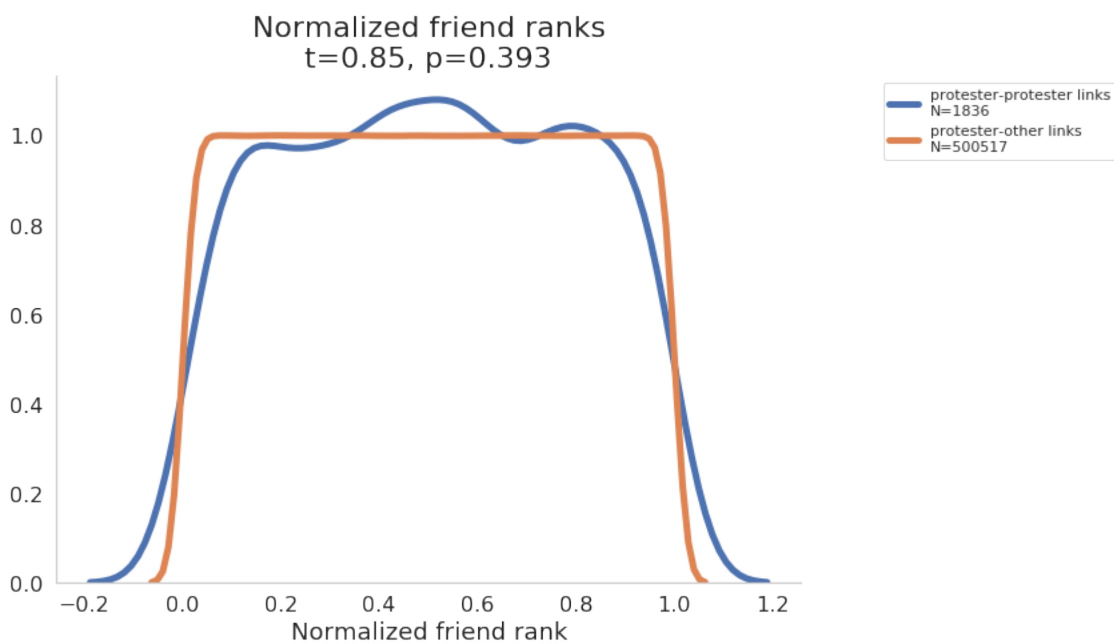

Figure 26: Comparing the rank of protester-protester ties to the rank of protester-nonprotester ties. The similarity in distributions suggests that protester-protester ties were not largely formed after the protest due to meeting and becoming friends at the protest.

2.9 Direct Evidence of Exposure on Twitter

For our characterization of social theories of protest to hold, a person must be exposed to her social ties' sentiment toward the protest. We measure social ties using following relationships on twitter, which we argue is a valid measure of underlying social networks. While interactions on Twitter are only one of many possible channels of exposure to a social tie, we can use the tweets that a user sends as a measure of the content to which followers are exposed on Twitter. This is an incomplete test– if none of person i's social contacts on Twitter post content about an upcoming protest, this does not mean that person i is not exposed to her social contacts' views toward the protest. i may be exposed to these ties on other social media platforms, or in real world interactions of various kinds, and so may learn of their sentiment toward the protest through other means. If i's social contacts *do* post content on Twitter about the protest, then we know we are observing a source of potential exposure.

We randomly sample 200 users from our Protesters set and manually classify the content of Tweets that the user shares in the days leading up to the protest. Specifically, for each user, we generate a page that looks like the Twitter feed of the user, and populate it with the tweets that user sent during the three days before the Charlie Hebdo protest. Figure 27 shows an example our replicated Twitter feeds. We then had two French-speaking undergraduate students read each page and classify its content on the crowd sourcing platform CrowdFlower. The coders were told that these were French accounts' content from the days between the Charlie Hebdo massacre and the protest, but were *not* told that these accounts sent Tweets geocoded to be present at the protest. We asked coders whether the account contained any tweets with certain content. When the coders disagree, we use the judgment of the coder with the highest confidence score on CrowdFlower. The comparisons all hold if we restrict attention to unanimous judgments.

Their coding reveals that 44% of the sampled accounts tweeted specifically about the upcoming Charlie Hebdo protest (as opposed to merely tweeting about the issue of Charlie Hebdo, which was true of 100% of the accounts by construction of the sample). Of the accounts that tweeted specifically about the protest, 28% tweeted logistical information about the protest, 20% tweeted about their personal plan to attend the protest, and 13% tweeted an explicit encouragement of others to attend the protest. Given that twitter is far from a person's only source of exposure to her friends' intentions toward a protest, that we find direct evidence of opportunities for exposure in so many accounts on Twitter alone is reassuring that the theory uses a plausible mechanism.

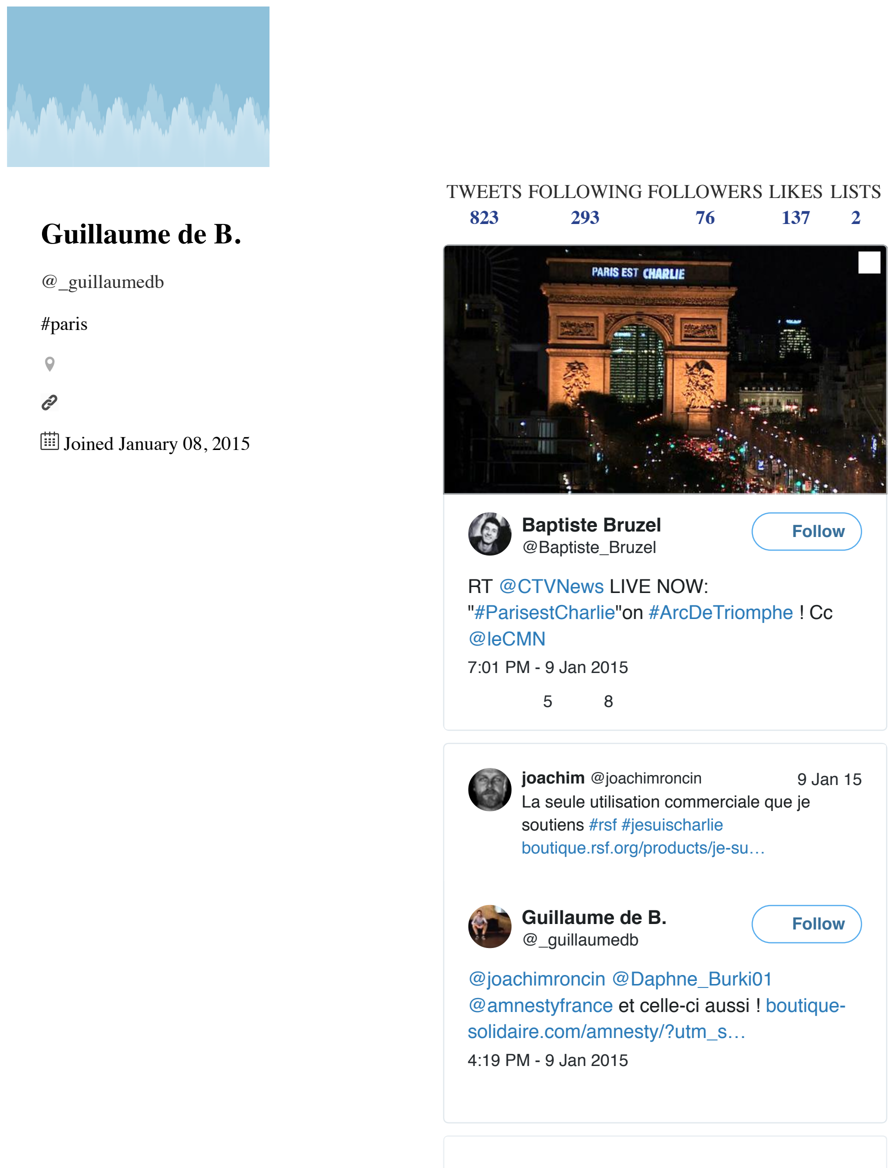

Figure 27: Example of the replica of a user's Twitter page generated from a user's Tweets between the massacre and the protest that coders saw when classifying Tweets.

Interestingly, when asked to make a guess about how likely it seemed that the user would eventually attend the protest (again, without knowing that all users in this sample did in fact attend), coders had the impression that more users might attend than the number of users who explicitly tweeted about the protest. Coders felt that 64% of this sample had at least an even chance of attending. Exposure may occur in subtle ways; a person may have influence on another's valuation of the protest even if she does not explicitly say "I value the protest highly." Simply by engaging with the issues surrounding the protest in a credible way, a person may signal to others the value of a protest.