Who Flees Conflict?

A Big-Data Approach to the

Determinants of Forced Migration *

Aidan Milliff † & Fotini Christia ‡

April 23, 2026

Abstract

Every year, millions of people encounter political violence and some leave their homes in pursuit of safety. What factors influence whether or not a person migrates to escape conflict? We use 63.5 million anonymized, geo-located cell phone records to study the migration behavior of over 60,000 people during a 2011-2012 conflict between the Yemeni Government and insurgents. We find that the structure of individuals' social and physical networks are important and nuanced predictors of migration. Specifically, we find that different types of social centrality have opposite marginal associations with probability of migration, and that accounting for social centrality moderates the connection between migration and important variables in the literature like economic status and exposure to violence. Our results suggest that different aspects of social networks bear on individual decision-making both by facilitating opportunities to leave contested territory, and by facilitating resilience-in-place to weather political violence.

Total Word Count: 7,352[1]

*Thank you to Paolo Bertolotti, Adam Kaplan, Michael Freedman, Chagai Weiss, four anonymous reviewers and audiences at Empirical Studies of Conflict 2022, PolNet 2024, and the MIT Security Studies Working Group for feedback. We appreciate HPC support from MIT Supercloud, and RA assistance from Manuel Quintero. This paper builds on results from a collaborative project with Michael Curry, Constantinos Daskalakis, Erik Demaine, John P. Dickerson, MohammadTaghi Hajiaghayi, Adam Hesterberg and Marina Knittel. We thank them all. The project is registered with the MIT Committee on the Use of Humans as Experimental Subjects both as protocol #1309005907 and exemption determination E-2274.

†Department of Political Science, Florida State University: amilliff@fsu.edu.

‡IDSS and Department of Political Science, Massachusetts Institute of Technology. Contact: cfotini@mit.edu.

[1]Body: 3,226; Methods (§5): 1,558; Tables, Figures, Captions: 980; References: 1,588

1 Introduction

How do people decide whether to flee their homes during conflict? People living in conflict zones often face circumstances in which both remaining and fleeing carry substantial, hard to measure risks. Previous empirical research finds support for different (and competing) explanations in different cases: more people leave areas where violence is more intense (Weiner, 1996; Davenport et al., 2003; Moore and Shellman, 2004; Czaika and Kis-Katos, 2009, among others), people with more (liquid) economic resources are more likely to flee (Schon, 2019). Better social resources are sometimes associated with remaining (Adhikari, 2013; Marston, 2020) and sometimes with leaving (Schon, 2020). More satisfaction with governance predicts remaining (Revkin, 2019), and higher individual risk tolerance may either promote staying (Mironova et al., 2019) or migrating (Camarena et al., 2020).[2]

We use over 63.5 million anonymized, time- and location-stamped Call Detail Records (CDRs, often called cell phone metadata), to study the migration behavior of over 60,000 cell subscribers in Yemen's Abyan governorate as they navigate violence and fluid territorial control in 2011 and 2012. Compared to typical survey- and interview-based data collection, CDRs provide important advantages for studying forced displacement. They are collected in real time and without self-reporting, which avoids the threats of faulty memories and post-treatment interference which are common challenges for migration studies in social science. The data also provide information—like months of call history and location history, or precise measurement of social networks—that no reasonable individual *could* report in a survey.

Our analyses show that different aspects of social network structure—beyond the "strength" or "weakness" of a person's network—have different associations with migration propensity in the Abyan conflict. We also find much smaller but still significant associations between migration decisions and pre-migration mobility, another factor under-explored in previous literature. Incorporating these newly-measured factors into traditional migration models changes the estimated effects of resource access and violence exposure, which have been the primary foci of social science research on forced migration

[2] We cite and discuss the existing literature more thoroughly in Appendix A.

Different aspects of an individual's mobility and social network affect migration during the Abyan conflict in different ways. First, we find large and opposite-signed effects for different measures of social centrality: people who occupy more "reachable" or influential positions (measured by PageRank) in pre-conflict communication networks are likelier to migrate during conflict while people who have overall broader networks (higher degree centrality) are more likely to remain. This finding sheds new light on the social dynamics of migration in a conflict setting. Showing that individuals in influential positions are likelier to migrate provides new insight into the possible causes of the "tipping point" phenomenon in refugee crises: When influential people migrate, they might simultaneously a) inspire others in their networks, b) communicate practical information about migration back to people in their home communities (Holland and Peters, 2020), and c) decrease the potential resilience of the community they leave behind. Second, we find that more mobile individuals *during* conflict and those with geographically wider call networks before conflict are slightly more likely to migrate. The mobility results are intuitive, but not previously documented. In this article, we introduce our approach to modeling migration decisions, describe the data sources, present results, and discuss the results' importance for understanding how social influence and mobility affect forced migration.

2 Yemen's War: 2010-2012

We study displacement in the 2011-2012 conflict between the Yemen government and Al-Qaeda in the Arabian Peninsula (AQAP). The broader conflict stretches back into the late 1990s, but first erupted into serious violence with a 2009 government offensive against the newly-unified AQAP (Uppsala Conflict Data Program, 2020). In March 2011, AQAP fighters captured major population centers in Abyan, a governorate in Southwestern Yemen just east of the port city of Aden (see Figure A.2), initiating a roughly 15 month period of sustained conflict and over 2,500 recorded battle deaths in a region which had previously seen little violence (See Figure 1d for a timeline, and Figure A.3 for geographic distribution).

Violence displaced a massive proportion of Abyan's pre-conflict population of 525,000 residents. In our data, the weekly displacement rate reaches as high as 70% at the height of the government offensive in 2012, but is not strictly increasing over time: Some people move in and out of

Abyan repeatedly during the conflict (see Figure A.1).[3] Consistent with contemporaneous reporting (UNOCHA, 2013), our data show that large numbers of displaced people land in Aden, only about 65 kilometers away from Zinjibar, Abyan's capital (Compare Figure 1a to Figure 1b).

3 Using Telecom Records to Study Behavior during Conflict

Researchers are increasingly using individual-level cell data to study political violence (Salah et al., 2019; Freedman et al., 2021; Aydoğdu et al., 2025), since communication technology is widely recognized as a tool for measuring conflict behavior (Christia et al., 2015) and as a factor influencing it (Shapiro and Weidmann, 2013; Pierskalla and Hollenbach, 2013; Walter, 2017; Christensen and Garfias, 2018). Meanwhile, a growing body of research in economics and computer science has developed methods to *enhance* the limited attributes available in raw communications data to study social and political phenomena (See Appendix C). For example, Blumenstock et al. (2015), Pappalardo et al. (2016), and Felbo et al. (2017) predict individual economic status from CDRs and validate these predictions against direct wealth measurements. Similarly, Bozcaga et al. (2019) constructs a religiosity measure tailored to users in Muslim-majority countries, while Christia et al. (2021) uses location and mobility to estimates attitudes toward combatants in a civil conflict. In our study, we apply some of these techniques to augment CDR data for modeling migration, with a detailed explanation of this process provided below in Section 5 and Appendix D.

Cell phone data provides a good window into the behavior of conflict-affected Yemenis for two reasons. First, CDR data represent a major improvement in the quality and scale at which we can measure (and model the effects of) individual communication behavior, social networks, and mobility. Second, qualitative evidence from the study of displacement suggests that cell connectivity is extremely important in the daily lives of civilians affected by conflict (mobile internet was very

[3]The UN Office for the Coordination of Humanitarian Affairs (UNOCHA) registered over 236,000 IDPs from Abyan, but their statistics likely undercount total displacement since they only track individuals who engage with their agency for services, rather than migrating to stay with family or acquaintances. Comparing official refugee counts to refugee cell activity in Turkey, for instance, the Data for Refugees project finds evidence of under-counting by the UN.

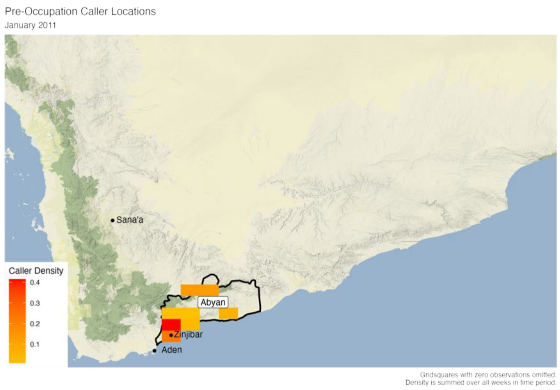

(a) Subscriber call locations in January 2011, the first month in which baseline data are collected. Abyan Governorate boundaries outlined in black.

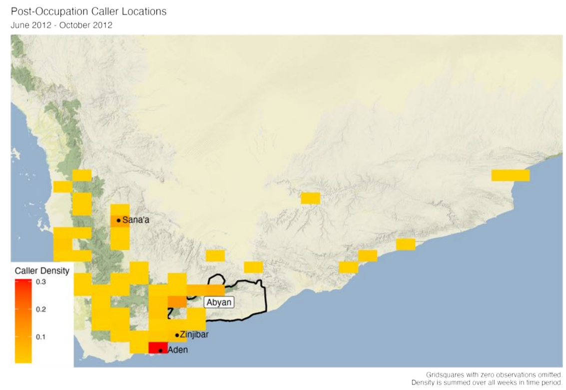

(b) Subscriber call locations after the end of the occupation, from June 2012- October 2012. Abyan Governorate boundaries outlined in black.

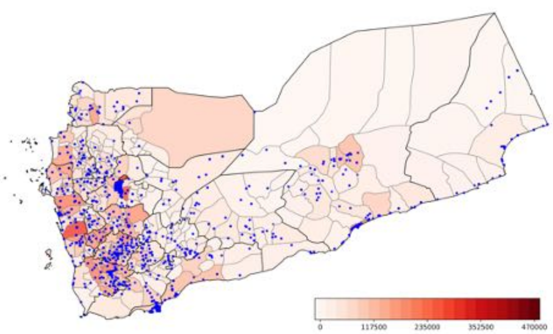

(c) Tower locations in Yemen, overlaid on district population density.

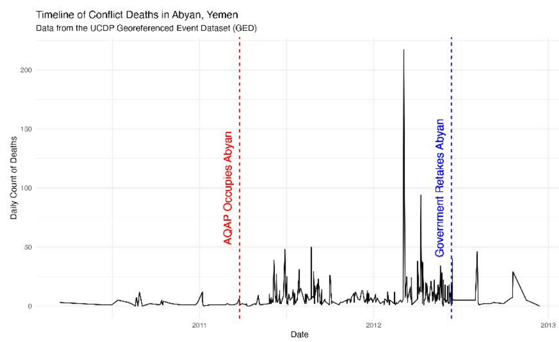

(d) Timeline of UCDP fatalities in Abyan, showing a sharp spike in "conflict" that begins with the AQAP occupation in March 2011 and persists through the government counteroffensive in early 2012.

Figure 1: Heatmap of Abyan residents' locations. Abyan residents were identified via their calling behavior in the baseline period of January 2011. Only subscribers whose "home locations" were in Abyan, are tracked as residents in our study. After 15 months of conflict, pre-conflict residents of Abyan are dispersed across the country, but many remain in Abyan. The third panel shows tower locations, the locations to which our method is able to resolve "home locations." 2G towers, which have an ideal-condition coverage radius of 35km, can capture presence in all populated districts of Yemen. The fourth panel shows a timeline of violent events in Abyan (measured by fatalities), indicating that the start of the occupation marks a major shift from "non-conflict" to "conflict' 'in the Governorate.

rare in Yemen from 2010-2012). Previous studies in the Middle East find that cell phones are the main channel that displaced people use to communicate both with relatives who have not migrated (AFAD, 2013) and with humanitarian organizations (Al-Sabahi and De Santis, 2016), and that access to cell phones is especially a high priority for the many migrants who are separated from their families during displacement (IOM, 2021). Of course, CDR data also come with limitations, which we discuss in detail in Appendix D. In brief, the results below should be interpreted with the understanding that: 1) CDR data register users' locations when they make calls or texts, not consistently like newer GPS data; 2) Cell coverage is country specific, making it difficult to measure international migration if/when it occurs; and 3) The correlation between individual and cellphone line is not measurable using CDR data, and is presumably not always 1:1.

4 Generating New Hypotheses on Social Networks and Migration

The actual structure of social/communication networks, as opposed to self-reported social ties which are subject to well-known recall biases (Brewer, 2000; Marin, 2004), has been studied relatively little as a determinant of forced migration during conflict.[4]. CDR data measures networks of communication that are almost certainly more complete and reliable than what Larson (2021) calls "proof of concept" studies eliciting self-reported networks through surveys or interviews.

To develop and test new hypotheses about the effects of network structure on conflict displacement, we use the graph structure of our data to calculate a variety of social network metrics which disaggregate different types of network "strength" (See Section 5.2 and Appendix E). The 63.5 million records in our dataset include all calls and texts involving identified Abyan residents between January 2011 and October 2012. The pre-conflict social network, however, is constructed from only the pre-occupation subset of approximately 5.5 million records (January–March 2011). Because many pairs of subscribers exchange multiple calls, these records collapse to 760,819 unique undirected edges connecting 230,565 nodes—the graph on which we compute centrality statistics. We describe its construction in detail in Section 5.2; here we introduce the two centrality concepts at the core of our hypotheses.

We focus on two distinct concepts of network strength: breadth and influence. First, we measure network breadth, whether an individual "knows everyone" or has "popularity" in a community, by calculating degree centrality—the total number of connections an individual has, normalized by the number of connections that are possible given the size/structure of the network. We suggest, in

[4]The existence of social ties, however, is known to be a key predictor of economic migration flows and timing (Moretto and Vergalli, 2008; Holland and Peters, 2020), and an important cause of refugee return *after* forced migration (Arias et al., 2014; Alrababah et al., 2023).

the strength of weak ties spirit, that wider networks might support resilience in place (Finkel, 2017; Marston, 2020), or may be an artifact of local leadership roles which disincentivize fleeing. Recently, Blumenstock et al. (2023) find similar associations between degree centrality and *economic* migration.

Second, we measure individuals' influence within their social networks by calculating PageRank, which measures importance as a function of an individual's number of connections, and the importance/prominence of those connections (Page et al., 1998). PageRank, a widely-used measure of centrality, expands on the intuition of betweenness centrality (Padgett and Ansell, 1993). A node's PageRank depends on its degree, and the degree of every node with which it shares an edge, therefore "rewarding" nodes that are connected to nodes that are themselves influential/important.[5] We posit that influence will matter for migration behavior because individuals who occupy influential positions in a social network have, in effect, a social resource that they can leverage to facilitate out-migration during conflict. If socially-influential people are more able to pursue their chosen strategy (migrating or not) and if the desire to migrate is widespread due to environmental factors like high-intensity violence (see Schon, 2020 for this dynamic in Syria), we would expect to see a positive association between social influence and migration, *even if* some influential individuals are leveraging their status to resist migrating.

[5]We calculate both directed and un-directed PageRank (the latter does not distinguish between incoming and outgoing connections). Previous studies in the United States suggest that *receiving* calls is indicative of higher social status within a relationship (Dong et al., 2015), which supports the use of Directed Pagerank. However, in contexts like Yemen where cellphone users pre-pay for airtime, the social norms around call direction are often the opposite: For a call of any meaningful length, recipients are typically the party with no airtime to pay (often lower status), rather than the socially important party (Donner, 2007). We therefore use un-directed PageRank in our main models. Additional results in I.4 replicate our main specifications substituting *directed* PageRank. Results there support the same conclusions about PageRank centrality.

5 Methods: Measuring Human Behavior in CDRs

We augment raw CDRs to generate data about individuals' migration status, exposure to violence, social milieu, socioeconomic status, religious practice, and surroundings in Abyan. CDRs are provided with anonymous, unique, consistent subscriber IDs that allow us to follow particular subscribers over time.[6] In addition to timestamp, duration, and subscriber IDs for both caller and recipient (or texter and recipient), CDRs include tower IDs for both parties, which can be matched with tower location data to identify coarse caller/callee locations at the time of the communication. Compared to the modern GPS data increasingly used in human mobility research to identify precise device locations at regular intervals(e.g. Kraemer et al., 2020; Pappalardo et al., 2023), CDR-based location measures only guarantee that a subscriber is within the connection radius of a given tower (up to 35km in perfect conditions), and only at the time a call or text occurs. [7]

5.1 Identifying Abyan Residents and Measuring Home Locations

Our analyses focus on residents of Abyan as of January 2011.[8] For callers who connect to a tower in Abyan at least once in January 2011, we calculate four weekly "home locations"—which we define as the tower from which they make the greatest number of calls/texts in a week (the modal tower).[9] We discard subscribers whose home locations are outside of Abyan in any week January, and use the remaining 60,440 subscribers as our sample of pre-conflict Abyan residents. Assuming a rough cor-

[6] See Appendix G. To maintain anonymity, we never present data at the individual level.

[7] In practice, a subscriber's location is usually substantially less than 35km from the tower to which they connect. See Appendix H for more detail.

[8] We implicitly exclude any Abyan residents who are elsewhere during January 2011. This ensures that subscribers identified as "displaced" during conflict are exhibiting a substantial deviation from pre-conflict mobility/behavior.

[9] Most frequent location is one of many home detection rules used in the computational science literature (Pappalardo et al., 2021; Barreras and Watts, 2024). Results of the main subsequent analyses (time-series models) are robust to two alternative home detection rules from Pappalardo et al. (2021)–the tower to which residents connect on the greatest number of distinct days, and the modal tower on weekend days only. See Appendix I.6 for supplementary results.

respondence of one phone line to one person, our sample includes over 10% of the total pre-conflict population of Abyan, but given that cellphones are often shared between members of a household, our sample could include members of over half the households in Abyan.[10]

5.2 Measuring Social Networks

We calculate many different measures of *social centrality* by creating a graph from all the calls/texts that Abyan residents in the data make in January – March 2011, the three month window before AQAP occupation of Abyan.[11] The resulting graph has 230,565 nodes (60,440 Abyan residents and 170,125 non-resident contacts), and 760,819 un-directed edges. The graph has a total of 1,068 connected components, with the largest component including 227,568 nodes; 98.7% of all nodes in the graph. Using the entire graph, we calculate a variety of network centrality measures including PageRank, Degree Centrality, total degree, and clustering coefficient, described in greater detail in Table 2 and Appendix E. Though PageRank is a common and intuitive measure for the "influence" concept we seek to test, previous work has noted inconvenient properties including dependence on network structure (Ghoshal and Barabási, 2011). We show in Appendix I our results are robust to using alternative measures of direct influence including Harmonic centrality (Marchiori and Latora, 2000; Rochat, 2009) and Strength centrality (Barrat et al., 2004; Opsahl et al., 2010).

Centrality measures are calculated once per subscriber, using a time-limited graph of 5.5 million *pre-conflict* calls and texts that can include social alters from anywhere in Yemen. We restrict our focus to pre-conflict social centrality because the evolution of an individuals' communication network over the course of conflict is clearly endogenous to their migration behavior. We do measure other attributes of a subscribers' behavior—including percent pareto interactions[12] and call count—on a

[10] The average household size in rural Yemen (which includes Abyan) was estimated at 6.7 shortly after the time period of our study (https://dhsprogram.com/pubs/pdf/fr296/fr296.pdf).

[11] Communication research typically considers calls as signifying stronger connections than texts (Sapiezynski et al., 2018). We combine the two because texts are vanishingly rare in Yemen in 2011, making up 0.4% of records. Only 0.2% of edges in the combined call/text graph are formed solely by texts.

[12] Percentage of user's contacts accounting for 80% of interactions (de Montjoye et al., 2016).

weekly basis throughout the conflict, using the remaining 58 million CDRs.

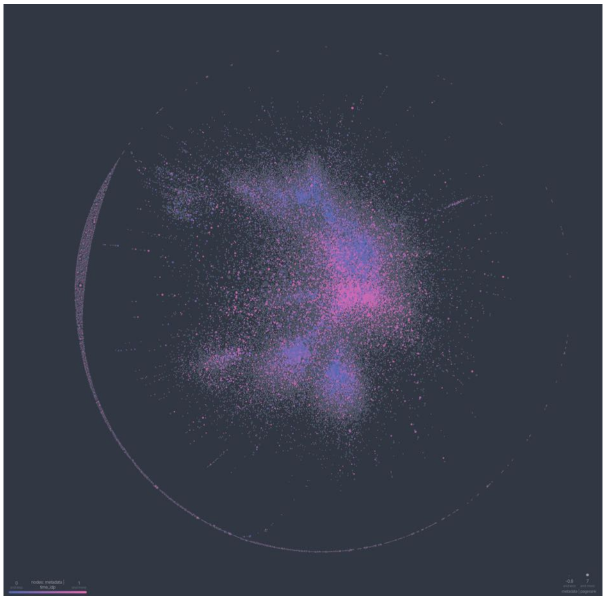

Figure 2: Static image of an un-directed projection of the graph of Abyan residents and their contacts, Jan–Mar 2011, created using `cosmograph` (Rokotyan et al., 2022). Brighter (pink) nodes are displaced for more weeks than darker (purple) nodes; grey nodes are non-resident contacts of ≥ 1 Abyan resident. An interactive visualization with the option to show directed edges is at bit.ly/abyan-callers.

Abyan Resident Callers: January - March 2011	
Nodes	230,565
Directed Edges	948,631
Undirected Edges	760,819
Connected Components	1,068
Largest Connected Component	227,568 nodes

Table 1: Descriptive statistics for Call/Text Graph including all identified Abyan residents (60,440), and their alters in other Yemen Governorates (170,125).

5.3 Measuring Migration and Mobility

We measure displacement as a binary variable, calculated once per subscriber-week, indicating whether or not the subscribers' weekly home location is *inside* or *outside* Abyan governorate. This is a conservative proxy for displacement, because crossing an administrative boundary is not necessary, for a person to meet the standard definition of internal displacement. This conservative measurement is appropriate given violence dynamics in Abyan in 2011-2012, which caused most displaced persons to leave for a different governorate like Lahij or Aden where territorial control was not contested or changing. Setting a high bar for "displacement" also gives reasonable assurance that we are not over-counting by capturing people who are simply living mobile lives in Abyan.

Some small number of individuals we identify as "displaced" are likely just traveling and would not consider themselves to have fled from violence. We argue that a number of other choices in measurement, combined with contextual features specific to Yemen, mitigate potential measurement error and minimize the threat that our findings result from a biased measure of displacement. First, all travel outside of Abyan after March 2011 (which counts as "displacement" if the travel duration is long enough that a subscriber's modal tower in a given week is outside Abyan) is implicitly compared to a baseline of no regular travel before the conflict, because any cell subscriber whose weekly home location is outside Abyan *at any point* in January 2011 is excluded from our population of study. Second, all models of displacement in the main text and Appendix I include measures of mobility (described below in this section) that capture subscribers' weekly average movement before the conflict (Radius of Gyration, pre-occupation, avg), and subscribers degree of movement in the prior week of the time series model (Radius of Gyration, person, lag). This means that the marginal associations between network statistics and displacement we describe can be interpreted as holding constant a subscriber's weekly movement, both in the baseline period and in the prior period of the time series. We expect that these adjustments mitigate the threat of bias in our estimated network–displacement associations.

5.4 Measuring Covariates

We construct a range of covariates to facilitate clearer interpretation of the marginal associations between social centrality, mobility, and displacement. The constructed covariates listed in Table 2 include summary statistics describing calling behavior, characterizations of non-displacement mobility, other social network attributes, and local environmental characteristics associated with a user's location.

First, we rely on results in the CDR literature and remote sensing literature (e.g. Blumenstock et al., 2015; Felbo et al., 2017, and others described in Appendix E) to construct proxies for subscribers' economic status. Previous work has documented associations between wealth and total call duration, propensity to initiate vs. receive calls, and clustering coefficient. We calculate each of these measures. Second, we calculate additional communication and mobility statistics for each subscriber as covariates. We calculate subscribers' weekly number of contacts,[13] the concentration

[13]This is distinct from their degree, and conceptually measures the proportion of a subscriber's total connections that are active in an average week.

of a subscriber's communication within a sub-group of those contacts, the average distance between caller and callee, weekly mobility, and number of calls/texts. Third, we use tools from the CDR literature to estimate subscribers' religiosity. This measure is important in a conflict where one actor (AQAP) holds an extremist religious ideology, because more/less observant individuals might feel less/more pressure to leave territory under their control. We count the number of calls made in the hour before Friday Dhuhr (noon) prayers as a proxy for observance (Bozcaga et al., 2019).[14]

Finally, we combine individual-level mobility and communication data with environmental covariates that can be spatially linked to specific cell towers. These include geo-referenced conflict events from the Uppsala Conflict Data Program (Pettersson and Oberg, 2020) and remote-sensing estimates of local economic activity from the Aid Data Project (Goodman et al., 2019). We assume that individuals are "exposed" to the conflict or economic conditions present within the coverage area of the cell tower to which they are connected. After combining data sources, the resulting dataset measures migration behavior (and most predictors besides social network statistics) at the subscriber-week level, resolving environmental covariates based on the subscriber's "home location" or most-used tower in a given week. Table 2 summarizes the constructed variables.

6 Modeling the Determinants of Forced Migration in Abyan

We use two different routines to model displacement from Abyan (equations in Appendix I). Covariates in each model of displacement are measured at different levels of analysis, described in Table 3. First, we fit a time-invariant linear model with tower and district random effects. The dependent variable here is the proportion of all weeks in the conflict in which a subscriber has a home location outside Abyan. We assume that, compared to a baseline behavior of consistent presence in Abyan during January 2011, a subscriber who spends many weeks outside of Abyan during the conflict is functionally displaced, and further, that the probability they are displaced rather than traveling increases with the length of time they spend outside of Abyan.[15] This serves as the best comparison to existing individual-level studies, which mostly model conflict migration as a "one-shot" process.

[14] Appendix I shows our results are also robust to measures based on pre-dawn (Fajr) and sunset (Maghrib) prayers, which are less widely observed and thus a higher bar for religiosity.

[15] We discuss measurement further in Appendix E.

Variable	Description	Citation/Source/ `python lib`
Call Duration	Grand mean of weekly average duration of calls.	`bandicoot` by de Montjoye et al. (2016)
% Initiated Interactions	Weekly average percentage of calls initiated by user.	`bandicoot`
# of Contacts	The number of contacts a user interacted with in a given week.	`bandicoot`
Call Distance	Average distance between caller and recipient towers (km).	Author's calculation
Friday Prayer Calls	Count of calls made in the hour before Friday noon prayers, over entire pre-occupation window.	Author's calculation & Bozcaga et al. (2019)
Radius of Gyration	Measure of mobility: Equivalent distance of mass from center of gravity of weekly places visited.	`bandicoot` & Gonzalez et al. (2008)
PageRank	Scaled measure of "importance" in a network: A node's directed PageRank is based on the PageRank of the nodes pointing toward it, i.e. it measures a node's importance through it's neighbors, and the independent importance of those neighbors.	`networkx` & Page et al. (1998)
Clustering Coef	Measures density of a node's ties, or the degree to which a node's network approximates a "clique" in which each node shares an edge with every other node. Calculated as a ratio of existing edges to possible edges in a local graph.	`networkx` & Watts and Strogatz (1998)
Degree Centrality	The fraction of nodes in a graph that a given node is connected to, normalized by the number of nodes in a graph.	`networkx`
Latitude/Longitude	Coordinates of the *tower* to which user connects.	Author's Data
Log GDP	Economic activity, measured at the 1km gridsquare level using DSMP nighttime light emissions and LandScan populationd data. We use the natural log of GDP associated with the km gridsquare that contains each tower.	Ghosh et al. (2010) *via* Goodman et al. (2019)
Pre-2010 GED Events	Sum of UCDP Georeferenced Events Dataset violence events that occur from 1989-2010 within 10 km of a given location (a tower).	Uppsala Conflict Data Program (2020) *via* Goodman et al. (2019)
GED Event Count	Count of UCDP GED events that occur within 10km of a tower in a given week in 2011-2012.	Uppsala Conflict Data Program (2020) *via* Goodman et al. (2019)
# Records	Number of calls/texts a user conducts in a given week.	Author's calculation
% Pareto Interactions	Proportion of user's contacts making up 80% of their interactions.	`bandicoot`

Table 2: Variable descriptions and sources. Some variables are calculated separately across multiple time periods; see Table 3 for the time periods over which variables are calculated for different regressions.

Results in Figure 3a shows that rarely-available network metrics substantially alter canonical associations between migration and predictors like wealth and violence. Coefficients in Figure 3a estimate the association of standard deviation change in the predictor variable with the number of conflict weeks outside Abyan.

The strongest predictors of migration are social: Respondents who have wider pre-conflict circles of active communication and higher numbers of distinct contacts in a week, are more likely to be outside of Abyan during the conflict. Pre-conflict PageRank and mobility (radius of gyration) slightly predict migration, but degree centrality is strongly negatively associated with migration.[16] The coefficient for local GDP where the subscriber resides (a proxy for subscriber wealth which is a canonically important predictor of migration) is insignificantly associated with time outside of Abyan. However, two other individual-level measures of poverty/wealth—percent of calls initiated vs. received and network clustering coefficient (See Appendix E)—have small, significant associations with migration, with the coefficient signs in line with predictions from the literature. Including network metrics not only changes coefficients, but also explains more variation. Appendix I, Table A.10 shows that the the fully-saturated model has a marginal R^2 that is 1.78 times higher than an otherwise similar specification without measures of PageRank, degree centrality, and clustering coefficient.

Second, we fit a time-series model to account for over-time variation in subscribers' location, behavior, and environment. Here, the dependent variable is a binary indicator for whether the subscriber is inside or outside of Abyan in a given week. We specify a logit model at the subscriber-week level with subscriber and week random effects and use Stan's Hamiltonian Monte Carlo sampler for estimation (See Appendix J for information and diagnostics). Figure 3b shows average partial effect estimates (King et al., 2000; Imai et al., 2008) from the random effects logit model, i.e. the marginal change in probability of being displaced associated with a shift from the 25th- to 75th-percentile value of a given predictor. Appendix I shows raw parameter estimates.

Time-series model results differ some but not all ways. First, the associations of most network statistics (calculated based on a snapshot of pre-occupation calling behavior) persist. Degree centrality is substantially negatively associated with migration, and PageRank is substantially positively associated (again, see discussion with Table A.11). The 25th to 75th percentile difference in each measure is associated with roughly a 15% lower probability of being outside Abyan and a 4% higher probability, respectively. Interestingly, the weekly average number of unique con-

[16]Centrality measures are empirically correlated. Table A.11 compares models omitting each. Also see Figure A.5.

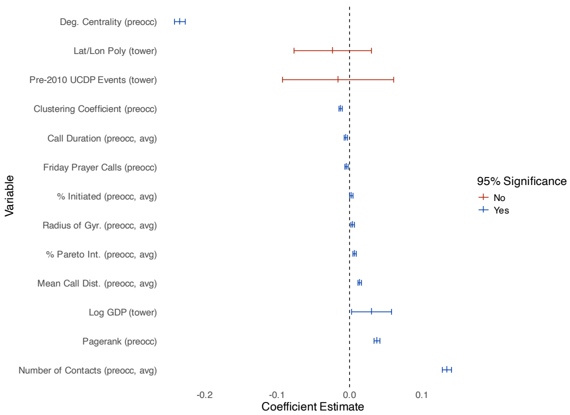

(a) Coefficient estimates (95% CIs) from a time-invariant linear model with district and tower random effects. Data are scaled and demeaned, so estimates represent change in total time spent outside Abyan associated with a one standard-deviation change in a predictor.

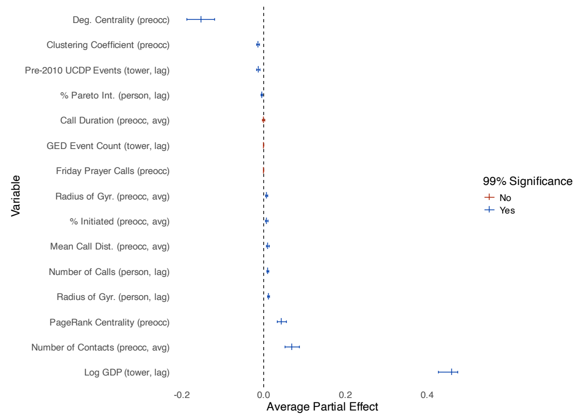

(b) Average partial effect of shifting from the 25th- to 75th-percentile value of the predictor on weekly probability of "fleeing" in the time-series model including subscriber and week fixed effects, as well as weekly calling behavior and environmental covariates. See Appendix Table A.3 for raw coefficients for these predictors. Error bars show 99% credible intervals from `stan`.

Figure 3: Model results newly identify social centrality and pre-displacement mobility as predictors of displacement behavior during conflict. See Appendix I for more specifications, and different measures of religious practice and social centrality

Time Period	Description
"Preocc, avg"	Mean of weekly measures between 1 Jan 2011 and 25 March 2011
"Preocc"	Single statistic calculated from CDRs 1 Jan 2011 - 25 March 2011
"Tower, lag"	Calculated for user's weekly modal tower ("home location") in a given week from 26 March 2011 - June 2012. Lagged by one time period vs. migration status.
"Person, lag"	Calculated from a user's weekly call behavior in a given week from 26 March 2011 - June 2012. Lagged by one time period vs. migration status.

Table 3: Time Period Descriptions for variables in all regression specifications.

tacts pre-occupation is substantially associated with migration: While the measures are positively correlated, having a wide network (degree centrality), is different from keeping up with that network (number of contacts). Once over time variation is modeled, we recover canonical associations between wealth and migration—local GDP shows a large positive association, and clustering coefficient (a proxy measure of poverty) is negatively associated.[17] We also see significant but sub-

[17]Because the GDP covariate is a one-week lag, the association with *leaving Abyan* may be overstated if people tend to flee to prosperous areas, and stay once they get there (See Appendix I.8).

stantively smaller associations between displacement and higher mobility behavior *during* conflict. Finally, the model shows no substantial association between a (short) lagged measure of localized violence and migration.[18] Geographic patterns in migration (latitude/longitude coefficients) indicate more out-migration from areas in southwestern Abyan, which were more violent. Violence appears to play a role in migration decisions, but longer term trends in violence, not short-term shifts at the margin, are apparently what matters.

Finally, we assess whether our results are specific to the conflict period rather than general features of mobility in Abyan. We use the pre-conflict period as a placebo (Bagrow et al., 2011), and model the migration behavior of subscribers present in Abyan during the first *two* weeks of January 2011 ($N \approx 52{,}000$) during the rest of the pre-conflict period (mid-January through March 2011). Placebo results indicate that our central findings are indeed specific to conflict: In the pre-conflict period, higher PageRank predicts *staying* in Abyan before the conflict rather than leaving (posterior mean $= +1.32$, 95% CI $[1.23, 1.43]$). The assoiation with degree centrality does not reverse, suggesting it captures stable local social embeddedness that retards mobility generally, rather than conflict-activated resource mobilization. Full placebo results are in Appendix I.7.

7 Discussion

We use previously-unavailable data sources to confirm some elements of existing theories of conflict-induced migration while also identifying important, previously un-observed predictors of migration behavior in Yemen. These results add insight into both the temporal and social dynamics of migration. We recover canonical associations between wealth and propensity to leave violent areas. We also find, intuitively, that high-violence areas are the places that cell subscribers are most likely to leave. After accounting for geographic location (highly correlated with long-term violence, in the Abyan conflict) weekly variation in violence has a vanishingly small association with migration propensity. We also find novel evidence of small "preparatory behaviors" that are leading indicators of a subscriber's appearance outside of Abyan. More mobility and making a greater number of calls overall are robustly associated with a slight increase in migration propensity. This may reflect

[18]Violence events are rare, so average partial effect estimates for reasonable ranges are precisely zero. The raw odds ratio (see Figure A.6) is significant, but minuscule compared to other associations.

people laying the groundwork to leave, perhaps reaching out to people outside of their geographic community to establish a "landing place."

Our key contribution is adding greater specificity and empirical leverage to research on social correlates of displacement, showing that people who are particularly influential or reachable as bridges/brokers in a network are more likely to leave, while people who, all else equal, have broadly-based social networks are more likely to remain.[19] Having a geographically-broader social network (measured by pre-conflict call distance) is slightly associated with migration as well. While intuitive—a person is increasingly likely to leave in the strength of their ties to distant places—this finding is novel in quantitative literature. These results parallel community-level social processes, e.g. reliance on social networks for information sharing, that other scholars including Larson and Lewis (2018); Larson (2021) have observed in nascent conflicts, but we show evidence that they matter for migration decisions as well. In line with this research, we argue that further examining social structure is a promising avenue of study for learning more about cascade dynamics in migration and other conflict-relevant behaviors (Kuran, 1989). Additionally, though it is not directly related to the causes of displacement we study in this paper, generating new empirical findings at the intersection of mobile communication and forced displacement using CDR data is useful because of the growing attention that the international humanitarian community is paying to mobile communications. The United Nations refugee agency (UNHCR), along with many governments and NGOs have identified mobile connectivity as an important pathway for "improv[ing] refugee wellbeing" (UNHCR, 2016). Our new findings about the association between mobile phone usage patterns and propensity to become displaced can be useful as part of this broader humanitarian initiative.

Some limitations of this study also warrant attention in future research. First, compared to GPS mobility data—largely unavailable pre-smartphones—our displacement measures are coarse, and

[19] This is notably different from the conclusions of recent research on intergroup conflict/cooperation by Atwell and Nathan (2022), which find that eigencentrality is associated with greater generosity in a public goods game. Our results are consistent with the idea that highly central individuals leverage that influence to leave their community behind.

potentially mis-measure displacement with error that would be observable in data collected today (See Appendix H for further detail). Second, our data offer a relatively "slim" view of behavior—and unlike GPS data only measures at all when subscribers use their phones—which limits our ability to make inferences about the intentions and perceptions of cell phone users as they move (or do not move) around Yemen. Correlating cell phone data with the attitudes of their users would be a major contribution in the future, though the logistical hurdles are high. Finally, our approach and data makes it difficult to measure *international* migration. New techniques in development (like Ellis, 2025), should make it possible to measure this behavior in the future.

Many questions about conflict migration remain unanswered. Our findings, in conversation with recent work on related topics like migrant destination, dynamics of refugee return, and "pull factors" add to that list : Why are different *types* of social status associated with different migration choices? What social signals—the information transmitted across social networks, which we do *not* measure—promote resilience and flight (Blumenstock et al., 2023)? To what degree, and by what social or informational mechanisms does the out-migration of influential individuals prompt cascading behavior by neighbors and contacts (Becker et al., 2021)? How do different network structures affect migrants' choice of destination (Steele, 2019; Munshi, 2020; Holland and Peters, 2020), and conversely, how do civilians choices about migration or remaining affect the structure of their social networks going forward? Finally, are the social mechanisms we identify unique to forced displacement during conflict, or do they travel to other circumstances like displacement due to natural disasters (Lu et al., 2012) or economic migration (Blumenstock et al., 2023)? Addressing these questions will, in many cases, require re-opening the black box of communication content—characterizing the information shared across mobile networks, which is *not* part of CDR data. Doing so will not only provide insight into critical decisions about migration, but also better knowledge for designing realistic, effective policy interventions that promote security, safety, and well-being for people caught in conflict.

Data Availability

The raw datasets analysed during the current study are not publicly available, per the terms of a data sharing agreement between the providing telecom company and [Author 2]. Replication *data* deposited at OSF includes the subscriber-week summaries used in regression analyses, with random noise added to prevent re-identification. Inquiries to [Author 2] are welcome. See Appendix G for more detail.

Code Availability

Code and model objects necessary to reproduce the analyses in this study are deposited in the Center for Open Sciences OSF repository: https://osf.io/gstmk.

References

Adhikari, P. (2013). Conflict-induced displacement, understanding the causes of flight. *American Journal of Political Science 57*(1), 82–89.

AFAD (2013). Syrian refugees in turkey, 2013 field survey results. Technical report, Republic of Turkey Prime Ministry.

Al-Sabahi, M. and F. A. De Santis (2016). The role of cultural norms and local power structures in Yemen. *Forced Migration Review 53*, 55–6.

Alrababah, A., D. Masterson, M. Casalis, D. Hangartner, and J. Weinstein (2023, October). The Dynamics of Refugee Return: Syrian Refugees and Their Migration Intentions. *British Journal of Political Science 53*(4), 1108–1131.

Arias, M. A., A. M. Ibáñez, and P. Querubin (2014). The desire to return during civil war: Evidence for internally displaced populations in colombia. *Peace Economics, Peace Science and Public Policy 20*(1), 209–233.

Atwell, P. and N. L. Nathan (2022). Channels for influence or maps of behavior? a field experiment on social networks and cooperation. *American Journal of Political Science 66*(3), 696–713.

Aydoğdu, B., O. Bilgili, S. Güneş, and A. A. Salah (2025, January). Mobile phone data for anticipating displacements: practices, opportunities, and challenges. *Data & Policy 7*, e5.

Bagrow, J. P., D. Wang, and A.-L. Barabási (2011). Collective response of human populations to large-scale emergencies. *PLoS ONE 6*(3), e17680.

Barrat, A., M. Barthélemy, R. Pastor-Satorras, and A. Vespignani (2004, March). The architecture of complex weighted networks. *Proceedings of the National Academy of Sciences of the United States of America 101*(11), 3747–3752.

Barreras, F. and D. J. Watts (2024, June). The exciting potential and daunting challenge of using GPS human-mobility data for epidemic modeling. *Nature Computational Science 4*(6), 398–411.

Becker, S. O., V. Lindenthal, S. Mukand, and F. Waldinger (2021). Persecution and Escape: Professional Networks and High-Skilled Emigration from Nazi Germany. SoDa Laboratories Working Paper Series 2021-02, Monash University.

Blumenstock, J., G. Cadamuro, and R. On (2015). Predicting poverty and wealth from mobile phone data. *Science 350*(6264), 1073–76.

Blumenstock, J., G. Chi, and X. Tan (2023). Migration and the value of social networks. *Review of Economic Studies*.

Bozcaga, T., F. Christia, E. Harwood, C. Daskalakis, and C. Papademetriou (2019). Syrian refugee integration in turkey: Evidence from call detail records. In A. A. Salah, A. Pentland, B. Lepri, and E. Letouzé (Eds.), *Guide to Mobile Data Analytics in Refugee Scenarios, The 'Data for Refugees Challenge' Study*, pp. 223–249. Springer.

Brewer, D. D. (2000). Forgetting in the recall-based elicitation of personal and social networks. *Social networks 22*(1), 29–43.

Camarena, K. R., S. Claudy, J. Wang, and A. L. Wright (2020, 07). Political and environmental risks influence migration and human smuggling across the mediterranean sea. *PLOS ONE 15*(7).

Christensen, D. and F. Garfias (2018). Can You Hear Me Now? How Communication Technology Affects Protest and Repression. *Quarterly journal of political science 13*(1), 89–117.

Christia, F., M. Curry, C. Daskalakis, E. Demaine, J. P. Dickerson, M. Hajiaghayi, A. Hesterberg, M. Knittel, and A. Milliff (2021). Scalable equilibrium computation in multi-agent influence games on networks. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI-21)*, Palo Alto. Association for the Advancement of Artificial Intelligence.

Christia, F., L. Yao, S. Wittels, and J. Leskovec (2015). Yemen calling: Seven things cell data reveal about life in the republic. *Foreign Affairs*.

Czaika, M. and K. Kis-Katos (2009). Civil conflict and displacement: Village-level determinants of forced migration in aceh. *Journal of Peace Research 46*(3), 399–418.

Davenport, C., W. H. Moore, and S. Poe (2003). Sometimes you just have to leave: Domestic threats and forced migration, 1964-1989. *International Interactions 29*(1), 27–55.

de Montjoye, Y.-A., L. Rocher, and A. S. Pentland (2016). bandicoot: a python toolbox for mobile phone metadata. *Journal of Machine Learning Research 17*(175), 1–5.

Dong, Y., J. Tang, N. V. Chawla, T. Lou, Y. Yang, and B. Wang (2015, March). Inferring Social Status and Rich Club Effects in Enterprise Communication Networks. *PLOS ONE 10*(3), e0119446.

Donner, J. (2007). The Rules of Beeping: Exchanging Messages Via Intentional "Missed Calls" on Mobile Phones. *Journal of Computer-Mediated Communication 13*(1), 1–22.

Ellis, R. (2025). International migration amid insurgent violence. Technical report, Georgia Institute of Technology, Atlanta, GA.

Felbo, B., P. Sundsøy, A. S. Pentland, S. Lehmann, and Y.-A. de Montjoye (2017). Modeling the temporal nature of human behavior for demographics prediction. Arxiv preprint, MIT Media Lab, Cambridge, MA.

Finkel, E. (2017). *Ordinary Jews: Choice and Survival during the Holocaust*. Princeton, NJ: Princeton University Press.

Freedman, M., F. Christia, S. Zoumpoulis, L. Yao, and A. Jadbabaei (2021). The effect of drone strikes on civilian communication: Evidence from yemen. *Political Science Research and Methods*.

Ghoshal, G. and A.-L. Barabási (2011, July). Ranking stability and super-stable nodes in complex networks. *Nature Communications 2*(1), 394. Publisher: Nature Publishing Group.

Goodman, S., A. BenYishay, M. Lv, and D. Runfola (2019). Geoquery: Integrating hpc systems and public web-based geospatial data tools. *Computers and Geosciences 122*, 103–112.

Holland, A. C. and M. E. Peters (2020). Explaining migration timing: Political information and opportunities. *International Organization 74*(3), 560–583.

Imai, K., G. King, and O. Lau (2008). Toward a common framework for statistical analysis and development. *Journal of Computational and Graphical Statistics 17*(4), 892–913.

IOM (2021). Iom yemen quarterly update, quarter 2. Technical report, International Organization for Migration.

King, G., M. Tomz, and J. Wittenberg (2000). Making the most of statistical analyses: Improving interpretation and presentation. *American Journal of Political Science 44*, 347–361.

Kraemer, M. U. G., A. Sadilek, Q. Zhang, N. A. Marchal, G. Tuli, E. L. Cohn, Y. Hswen, T. A. Perkins, D. L. Smith, R. C. Reiner, and J. S. Brownstein (2020, August). Mapping global variation in human mobility. *Nature Human Behaviour 4*(8), 800–810. Publisher: Nature Publishing Group.

Kuran, T. (1989). Sparks and prairie fires: A theory of unanticipated political revolution. *Public Choice 61*(1), 41–74.

Larson, J. M. (2021). Networks of conflict and cooperation. *Annual Review of Political Science 24*, 89–107.

Larson, J. M. and J. I. Lewis (2018). Rumors, kinship networks, and rebel group formation. *International Organization 72*(4), 871–903.

Lu, X., L. Bengtsson, and P. Holme (2012). Predictability of population displacement after the 2010 haiti earthquake. *Proceedings of the National Academy of Sciences 109*(29), 11576–11581.

Marchiori, M. and V. Latora (2000, October). Harmony in the small-world. *Physica A: Statistical Mechanics and its Applications 285*(3), 539–546.

Marin, A. (2004). Are respondents more likely to list alters with certain characteristics?: Implications for name generator data. *Social networks 26*(4), 289–307.

Marston, J. (2020). Resisting displacement: Leveraging interpersonal ties to remain despite criminal violence in medellin colombia. *Comparative Political Studies 53*(13), 1995–2028.

Mironova, V., L. Mrie, and S. Whitt (2019). Risk tolerance during conflict: Evidence from aleppo, syria. *Journal of Peace Research 56*(6), 767–782.

Moore, W. H. and S. M. Shellman (2004). Fear of persecution: Forced migration, 1952-1995. *Journal of Conflict Resolution 48*(5), 723–745.

Moretto, M. and S. Vergalli (2008). Migration dynamics. *Journal of Economics 93*(3), 223–265.

Munshi, K. (2020). Social networks and migration. *Annual Review of Economics 12*(1), 503–524.

Opsahl, T., F. Agneessens, and J. Skvoretz (2010, July). Node centrality in weighted networks: Generalizing degree and shortest paths. *Social Networks 32*(3), 245–251.

Padgett, J. F. and C. K. Ansell (1993). Robust Action and the Rise of the Medici, 1400-1434. *American Journal of Sociology 98*(6).

Page, L., S. Brin, R. Motwani, and T. Winograd (1998). The pagerank citation ranking: Bringing order to the web. Technical Report SIDL-WP-1999-0120, Stanford University InfoLab, Stanford, CA.

Pappalardo, L., L. Ferres, M. Sacasa, C. Cattuto, and L. Bravo (2021, December). Evaluation of home detection algorithms on mobile phone data using individual-level ground truth. *EPJ Data Science 10*(1), 1–19.

Pappalardo, L., E. Manley, V. Sekara, and L. Alessandretti (2023). Future directions in human mobility science. *Nature computational science 3*(7), 588–600.

Pappalardo, L., M. Vanhoof, L. Gabrielli, Z. Smoreda, D. Pedreschi, and F. Giannotti (2016, December). An analytical framework to nowcast well-being using mobile phone data. *International Journal of Data Science and Analytics 2*(1), 75–92.

Pettersson, T. and M. Oberg (2020). Organized violence, 1989-2019. *Journal of peace Research 57*(4), 597–613.

Pierskalla, J. H. and F. M. Hollenbach (2013). Technology and Collective Action: The Effect of Cell Phone Coverage on Political Violence in Africa. *American Political Science Review 107*(02), 207–224.

Revkin, M. (2019). To stay or to leave? displacement decisions during rebel governance. Working paper, Yale University, New Haven, CT.

Rochat, Y. (2009). Closeness Centrality Extended to Unconnected Graphs: the Harmonic Centrality Index.

Rokotyan, N., O. Stukova, K. D., and D. Ovsyannikov (2022). Cosmograph: Gpu-accelerated force graph layout and rendering [computer software].

Salah, A. A., A. Pentland, B. Lepri, E. Letouzé, Y.-A. de Montjoye, X. Dong, Ö. Dağdelen, and

P. Vinck (2019). Introduction to the data for refugees challenge on mobility of syrian refugees in turkey. In A. A. Salah, A. Pentland, B. Lepri, and E. Letouzé (Eds.), *Guide to Mobile Data Analytics in Refugee Scenarios: The 'data for Refugees Challenge' Study*, pp. 3–27. Cham: Springer International Publishing.

Sapiezynski, P., A. Stopczynski, D. K. Wind, J. Leskovec, and S. Lehmann (2018, November). Offline Behaviors of Online Friends. arXiv:1811.03153 [cs].

Schon, J. (2019). Motivation and opportunity for conflict-induced migration: An analysis of syrian migration timing. *Journal of Peace Research 56*(1), 12–27.

Schon, J. (2020). *Surviving the War in Syria: Survival Strategies in a Time of Conflict*. Cambridge: Cambridge University Press.

Shapiro, J. N. and N. B. Weidmann (2013). Is the Phone Mightier than the Sword? Cell Phones and Insurgent Violence in Iraq. *International Organization 36*(2).

Steele, A. (2019). Civilian resettlement patterns in civil war. *Journal of Peace Research 56*(1), 28–41.

UNHCR (2016). Connecting Refugees. Technical report, UN High Commission for Refugees, Geneva.

UNOCHA (2013). Yemen humanitarian response plan 2013. Technical report, United Nations Office for the Coordination of Humanitarian Affairs, New York.

Uppsala Conflict Data Program (2020). conflict encyclopedia. Technical report, Uppsala University.

Walter, B. (2017). The new new civil wars. *Annual Review of Political Science 20*, 469–86.

Weiner, M. (1996). Bad neighbors, bad neighborhoods: An inquiry into the causes of refugee flows. *International Security 21*(1), 5–42.

Appendices

A Summary of Existing Migration Literature

Previous studies of migration specifically during conflict have been carried out in conditions that are challenging both for research design and the practicalities of data collection. We summarize the findings of this literature in Table A.1.[20] Two issues are common across previous answers to the question of "who flees?" These issues may limit the strength of the conclusions that can be drawn from the impressive data collection efforts. First, an early wave of quantitative studies of forced migration—including theoretically generative studies by Davenport et al. (2003) and Moore and Shellman (2004)—primarily use time-series cross-national data to explain the causes of refugee *flows*. These important data are not necessarily appropriate for explaining differences in migration behavior at the micro-level, comparing settlement to settlement or neighbor to neighbor (Robinson, 1950).

A second wave of studies address this lacuna and focus on individual behavior using structured interviews and household surveys. These studies are clearly better oriented to identify individual- and community-level migration determinants. At the same time, though, many are unable for practical reasons to observe and measure the determinants of choice *before or during* critical decision points. Many studies of migration, for example, only measure the behaviors and attitudes of refugees, choosing for understandable reasons of accessibility to explain outcomes like migration timing, rather than the choice to migrate or not. Some studies that surmount massive obstacles to collect data from "remainers" during or immediately after violent conflicts face another inferential challenge. Survey responses, behavioral game outcomes, or interview testimony that refugees and remainers give during or immediately after conflict might vary systematically as a result of experiences that would be downstream consequences of a decision to migrate or remain.

[20] We do not include the growing literature on refugee return (see some cited in the body of the paper) because it focuses on a different set of causes than the initial migration literature.

Paper	Outcome	Key Cause	Method	Context	n
Weiner (1996)	Refugee flows	Intra-state conflict ↑	Case comparison	Cross-national	N/A
Schmeidl (1997)	Refugee flows	"Generalized violence" ↑	Pooled time series	Cross-national	N/A
Davenport et al. (2003)	Refugee flows	"Personal integrity threats" ↑	Pooled time series	Cross-national	N/A
Moore and Shellman (2004)	Refugee flows	Threat of violence ↑	Pooled time series	Cross-national	N/A
Engel and Ibáñez (2007)	Individual migration	Non-liquid assets ↓	1x survey	Colombia	363
Czaika and Kis-Katos (2009)	Community-level migration	Violence events ↑	2x Census Data	Indonesia (Aceh)	5,200
Adhikari (2013)	Individual migration	Social networks ↓	1x Survey	Nepal	1,424
Balcells and Steele (2016)	Community-level migration	Ideological align. w/ combatants ↓	1x Electoral returns	Spain	654
Jampaklay et al. (2017)	Individual migration	Violence events ↑	1x Survey	Thailand	1,009
Schon (2019)	Individual migration *timing*	Wasta (money/connections) ↑	Interviews	Turkey (Syria)	178
Mironova et al. (2019)	Individual migration	Risk preferences ↓	1x Behavioral game	Syria	232
Revkin (2019)	Individual migration	Governance satisfaction ↓	1x Survey	Iraq	1,458
Alrababa'h et al. (2020)	Individual return	Violence at home ↓	1x Survey	Lebanon (Syria)	3,003
Fearon and Shaver (2020)	Refugee flows	Conflict deaths ↑	pooled time series	Cross-national	N/A
Camarena et al. (2020)	Refugee flows	Violence ↑/Transit risk ↓	Event data	Italy (Libya)	307,000
Marston (2020)	Individual migration	Ties to Local Leaders ↓	Survey	Colombia	618
Becker et al. (2021)	Individual migration	Ties to earlier migrants ↑	Co-authorship Net.	Nazi Germany	1,129
Milliff and Christia	Individual migration	Social Centrality ↓/↑ Violence ↑, Economic status ↑	CDR data in time-series	Yemen	60,000 units 2.87M obs. 63.5m calls

Table A.1: Causes of Conflict-induced Migration in the Literature

B Yemen Maps and Background

B.1 Background and Scope Conditions

Certain characteristics of Yemen, and the conflict between AQAP and the Yemen government, may be important scope conditions for the results we find. First, the conflict in Yemen in this period roughly approximated what Kalyvas and Balcells (2010) call "symmetric non-conventional" warfare.[21] The displacement dynamics of other types of warfare like irregular warfare or high-intensity conflict between conventional military forces may differ from what we find here: the centrality of civilians and civilian-combatant interaction is higher in irregular conflict (Kalyvas, 2006), and often lower in conventional warfare. The conflict in 2011-2012, unlike more recent violence in Yemen, did not prominently feature conventional attacks that were aimed to encourage displacement. We expect that the dynamics of displacement would be substantially different in a conflict where at least one combatant party tried to strategically induce civilian flight (Greenhill, 2010; Zhukov, 2013; Balcells and Steele, 2016). The final scope condition that we expect to be particularly relevant for our study is Yemen's social homogeneity. Both Abyan and the primiary destinations for Abyan migrants are in areas of Yemen that are predominately populated by Sunni Muslims. Ethnic nor sectarian identity was a notable cleavage in the Abyan conflict. We expect that patterns of displacement in ethnic wars and other wars dominated by ascriptive identity cleavages would be systematically different, especially with regard to the effects of social network structure on migration.

[21] The Yemen government had substantially more sophisticated military power than AQAP at the time, but we would argue that the war was not an "irregular" war because AQAP managed to hold territory in Abyan somewhat durably.

B.2 Maps

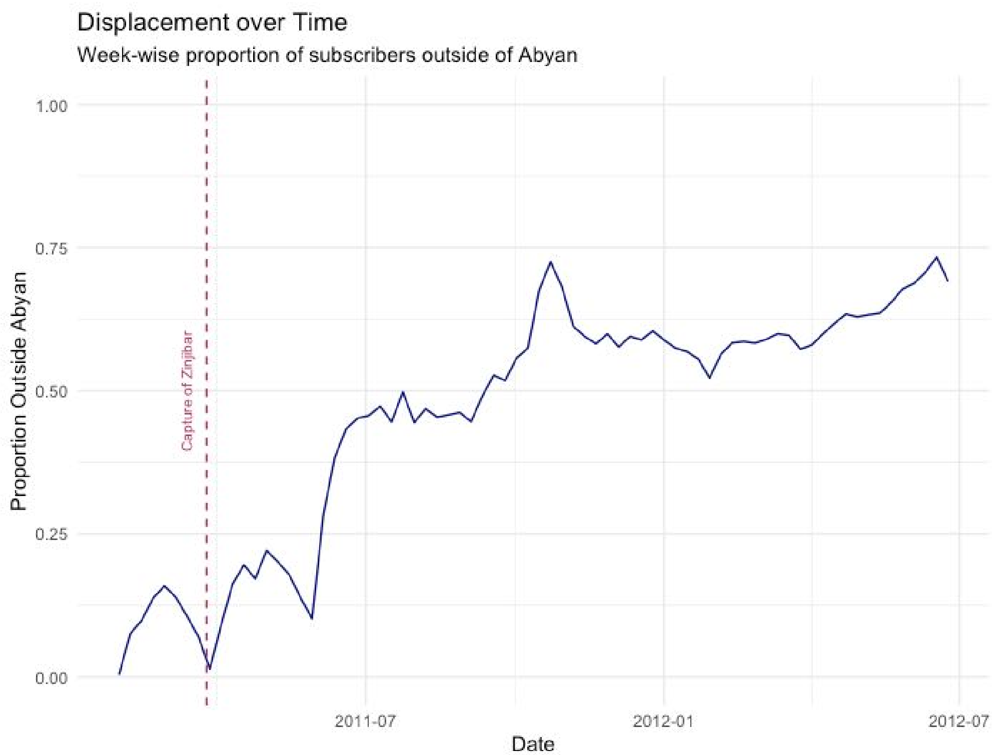

Figure A.1: Weekly proportion of subscribers outside Abyan. Vertical, dashed line indicates start of conflict.

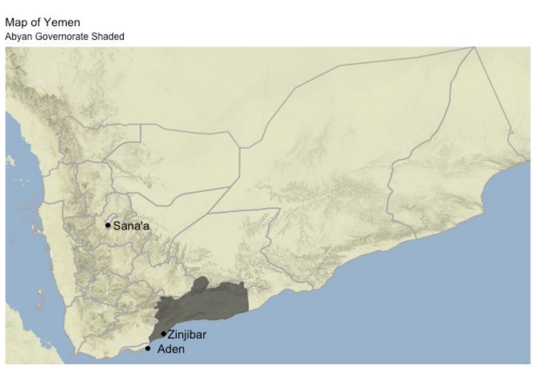

(a) A map of Yemen, with Abyan Governorate shaded. In addition to the capital Sana'a, and the port city of Aden, the map identifies Zinjibar, the capital of Abyan.

(b) A map of Abyan Governorate, identifying key population centers in and near the governorate, including the capital Zinjibar and the major port of Aden.

Figure A.2: Maps of Yemen and Abyan.

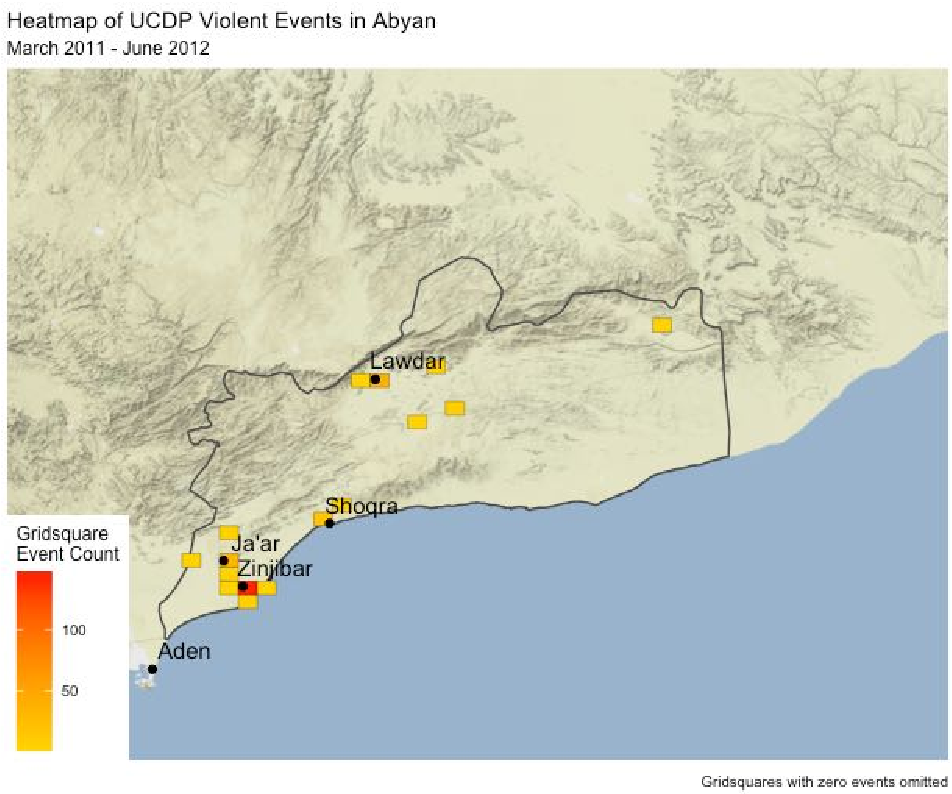

Figure A.3: Events in the UCDP Global Events Database for Abyan from March 2011 through June 2012. Violence is concentrated almost exclusively in settlements, with the greatest number of UCDP events occuring in and around Zinjibar, the capital of Abyan governorate. The map colors evenly-sized gridsquares according to the total count of separate events geolocated in the square. Square in which no events occurred are not plotted.

C Call Detail Record Literature

A substantial literature in computer science and economics use mobile communications data to study conflict relevant processes like the promulgation of ideas across networks (Bertolotti et al., 2024, 2020; Christia et al., 2021), poverty and mobility (Gonzalez et al., 2008; Blumenstock et al., 2010, 2015; Kreindler and Miyauchi, 2019), disaster response (Bagrow et al., 2011; Lu et al., 2012; Sîrbu et al., 2021; Yabe et al., 2022), non-conflict migration (Blumenstock et al., 2023; Aydoğdu et al., 2025), social integration (Salah et al., 2019), and the development/identification of social networks (Eagle et al., 2009; De Choudhury et al., 2010; Stopczynski et al., 2014).

One output of this literature is a set of new, empirically validated techniques for agumenting CDR data such that it can be used as a proxy to measure other social, behavioral, and economic attributes of interest, even identity features like age and gender in some case (Felbo et al., 2017). CDR data itself is quite a narrow measure of location and communication behavior (see Table A.2 for example rows from our data). Most of the analytical leverage researchers can get by using CDR data, therefore, come from computational techniques that rely on associations between calling/mobility behavior and other traits to infer those traits without direct measurement. We draw on some of these results to augment CDR data in our models of migration. We caution though, that some multivariate models linking CDR attributes to wealth, age, and gender, for instance, seem to be culturally specific—Blumenstock (2018) notes that models are "brittle" when applied to other social/cultural contexts. When using CDR attributes as proxies for wealth, we a) focus on attributes that the literature shows are consistently associated with wealth/income across a variety of contexts and b) often calculate multiple measures of a single underlying concept and look for them to be positively associated in our data (See Figure A.4). Because ground truth-validated age and gender prediction studies do not exist for most contexts outside Western Europe, and because we expect the association between cell activity and age/gender to be highly culturally specific, we do not attempt to measure age and gender using CDRs in our data. See Appendix E for more detail on the generalizable measures of economic status we use in this paper.

Time	**Anonymized IDs**		**Tower IDs**		**Duration**
Call Start	Caller	Recipient	Caller	Recipient	(seconds)
2012-06-13 7:57:12	34331363	30002125	3995	0623	1139
2012-06-13 7:57:13	20918574	38389599	1373	1290	10
2012-06-13 7:57:13	39496809	00146551	3338	1401	32

Table A.2: Three of $\sim$12 billion calls from the raw call detail records dataset, showing the data structure from which we construct a social network, as well as migration measures.

D Measurement Advantages and Limitations of CDR Data

Using Call Detail Records (CDRs) to measure and model forced displacement presents both opportunities and challenges compared to other common sources of data like interviews, surveys, and administrative records (see Appendix A). In terms of opportunities, CDR data addresses some of the major bias, attrition, and selection challenges that come with data sources like interviews and surveys (see brief discussion in the introduction). Many survey- and interview-based studies of migration decisions, for instance, can only reach people who have a) already chosen to migrate and b) ended up migrating to a destination with a high concentration of displaced people, like an IDP or refugee camp. This sampling problem—driven, to be sure, by genuine logistical and ethical challenges—creates a host of inferential issues ranging from selection on the dependent variable (migration status) to the possibility that selection into a study's sample might be correlated with the independent variable of interest (e.g. estimating the effect of wealth on migration probability using a sample of individuals in IDP camps might produce a biased estimate if wealthier migrants are likely to end up somewhere other than an IDP camp). Using CDRs mitigates many of these problems because the data source allows us to identify a population of conflict-affected individuals *ex ante* and track them over time as they migrate or remain. We argue in the body of the paper that this improvement in sampling strategy might account for some of the differences between the causes of migration we highlight and those that are most notable in the existing literature.

CDR data also provide a major advantage in terms of the quantity and quality of data that are available about each individual in the population. Many previous studies that look at social determinants of migration (economic migration and displacement during conflict) ask survey respondents or interviewees to report ties that they deem important in the context of their migration decision (e.g. whether they knew anyone in the destination to which they migrated). While people can probably report their social networks in a limited way (*pace* Nisbett and Wilson, 1977), there is certainly a limit to the completeness of those reports, and self-reports are logically likely to bias measurement of social networks toward two types of ties: recently-activated ties (i.e. social alters at the top of one's mind when the survey is conducted), or ties the respondent consciously believes are important to the question at hand. Evidence from early empirical network research, however, suggests that apparently important ties are *not* necessarily the most consequential ties in a network (Padgett and Ansell, 1993). CDRs, in contrast, provide a complete graph of communications over a long time period. The same is true with mobility. Interviewees or survey respondents might not be able to accurately report their whereabouts over a long time series—traumatic experiences like fleeing one's home may make it especially difficult to recall "peripheral" details like exact dates and times (Levine and Edelstein, 2010)—but CDR data facilitate measurement of location (and harder to self-report quantities like average weekly mobility) in a manner not prone to recall bias.

Along with these advantages, CDR data do have important limitations compared to other sources. We mention a number in the paper and describe them in more detail here. First, and most obvious, CDR data are "slim" insofar as they include very few features that might be relevant to predicting migration. Variables like social network status and calling behavior can be directly constructed from CDRs, but many other potentially-important data points for understanding migration decisions are either measured by proxy or wholly unobserved. Intentions and perceptions, which are important for many theories of migration cited above are fundamentally not possible to re-create from cell phone metadata. Other key variables including socioeconomic status, mobility, religiosity, etc. are measured using proxies based on observed correspondence in the economics and computer science literature about how people with certain traits use their phones. We describe why each proxy is reasonable and useful in the context of Yemen in section E below, and also describe why we decline to use certain proxies in the literature that have been found to be "brittle" or highly context dependent. All the same, we are not able to directly validate the proxies against "ground truth" data from Yemen.

Second, CDR data make it more challenging to 1) identify our population of interest (pre-conflict residents of Abyan) and 2) assign migration-status values to individuals than it would be with self-report data. In the paper, we infer residency status based on continuous presence in Abyan prior to the start of the conflict we study, but there are situations in which our measure has Type I and Type II errors. Type I error (individuals we mistakenly track through the conflict who are *not* residents of Abyan) could occur if individual cell subscribers make long-term visits to the governorate during our baseline residency-detection period, and then "happen" to return home during conflict. We believe this type of error is limited in the case of Abyan because the governorate simply is not a major *destination* for internal travellers from other parts of Yemen. As of the 2004 census (the last one before the 2011-2012 occupation) Abyan's two largest settlements had only tens of thousands of residents, and it's primary port city of Zinjibar is located only 50km from the vastly larger commercial capital of Aden. Type II error—mistakenly omitting residents of Abyan because they travel during the baseline period—is probably more common than Type I error, but in the particular context of this study, it may have the surprising effect of *reducing* measurement error in the dependent variable. Type II error here restricts our population of interest to individuals who do not frequently travel outside of Abyan for extended periods of time, which makes our assertion that staying outside Abyan for an extended period *during* conflict is evidence that a person is displaced, rather than simply away for other, temporary reasons.

A third limitation of CDR data in most cases is that, while the data for covered subscribers is extremely rich, the breadth of coverage has limits. In our case, there are three important limits to consider. First, our data come from one of three major cellular providers operating in Yemen at the time of the study, so we are missing data on most subscribers in the country. This would be a major threat to inference if there were evidence that subscribers to different cell networks were systematically different (e.g. one carrier was "lower cost," or brand new, or if different carriers had different technologies/coverage ranges as was the case in the United States for many years). There is not evidence in the context of Yemen that subscribers to one company or another would be systematically different; multiple major providers started around the same time, all seemed to have nation-wide customer bases. We believe, accordingly, that the subscribers we capture are a reasonable proxy for Yemeni cell subscribers broadly speaking. A second important limit to CDR data is that coverage does not extend beyond Yemen–if a person leaves Yemen their calls and texts stop appearing in our data. This is an important caveat for studies of displacement, since thousands of Yemenis did migrate abroad during the time period we study. In this case, the relative isolation of Abyan likely makes the data problem less severe. It would be difficult to leave Abyan for another country in 2011-2012 without first appearing in another part of Yemen that had either an international land border, a commercial international airport, or a *larger* commercial port. People who could manage travel out of a conflict zone and then onward out of Yemen *without* staying in-transit long enough that we record them as internally displaced are likely people who are more influential, wealthier, and with more experience traveling (i.e. more pre-conflict mobility). Since we find all three of these factors to have positive marginal associations with displacement, accounting for the types of internal migrants we are likely to be missing would most likely *increase* the magnitude of the results we report in the paper. Third, CDR data struggle to capture small, short duration movements, since they only record locations during calls/texts, and there is a well-known inverse correlation between in-the-moment mobility and the likelihood of placing or picking up a call (Zhao et al., 2016). We aggregate locations at the weekly level to deal with this issue.

The fourth limitation of CDR data in all settings, but especially in lower income countries, is that the cell phone line does not necessarily have a one-to-one correspondence with an individual person. We treat cell phones as proxies for their users in the paper, but there is some reason to think that cell phones operate more like shared household goods in the context we are studying. This limitation has three consequences for our results. First, while our data cover a number of cell phone lines equal to 10% of Abyan's census population, it may actually be covering the communication patterns of over 50% of Abyan households, which may modestly reduce concerns about the representativeness of our population. Second, *iff* cellphones are effectively household goods *and* displacement during conflict sometimes separates households (Chiovelli et al., 2021), we might expect that the phone goes with the household members who are traveling, and the household members who are remaining in place will borrow a phone from a friend or neighbor. This plausible pattern might also add to the apparent discrepancy between known-undercount UNOCHA numbers reported in the paper and the migration proportion we observe. We do not expect, though, that it would bias the results we present. Third, since households with a single cellphone are more likely to have male primary users (Barnes et al., 2025), our results are probably most accurate for characterizing the behavior of men.

E Constructed Variables

E.1 Violence

Political violence is a key characteristic of Abyan during our period of study, and exposure to violence is a prominent explanation for migration behavior. We use the Uppsala Conflict Data Program Global Events Dataset (GED, Pettersson and Oberg, 2020) as a source of geo-located violence data. UCDP data is widely used by social scientists and follows a rigorous multiple confirmation coding procedure, which requires that each event in the dataset is supported by multiple news articles. For the 400-plus GED events in Abyan in 2011-2012, we identify all cell towers within 10 miles of the GED event coordinates and label those towers as "affected" by the violent event in the week it occurred. GED event coordinates attach the event to the centroid of hte smallest geographic unit that is corroborated by published sources. If, for example, GED geolocates an event to the city of Ja'ar, we code all cell towers within 10 miles of the center of Ja'ar as "affected." If GED sources describe an event as occurring in Khanfir district but provide no more specific location, we code all cell towers within 10 miles of the centroid of Khanfir district as "affected." We identify over 41,000 tower-event-week triplets.

E.2 Economic Factors

We rely on proxies developed in the CDR literature as well as proxies from the remote sensing literature to measure economic status. Across a variety of contexts and different studies, Blumenstock et al. (2010, 2015) and Felbo et al. (2017) identify a few features as suitable proxies for relative wealth. We focus primarily on the proxies that are most closely connected to monetary resources and less on possible proxies that are mediated by social norms. We avoid multivariate wealth estimation/prediction models in favor of individual proxies for economic status, because Blumenstock (2018) finds the particular functional forms of the models are brittle and under-perform with different cultures of communication and telecom pricing structures. The individual predictors we use are associated with wealth across many contexts, so they should be more plausible than brittle multivariate models, especially in settings with low wealth inequality like Yemen. Wealth is positively associated with average call duration in countries (including Yemen) where pre-paid cell plans dominate. We measure pre-occupation average call duration to proxy wealth. In markets dominated by pre-paid phone plans (like Yemen), propensity to initiate calls—thus paying for the communication which is free to the recipient of the call—is a predictor of wealth. This social norm is strong enough in many contexts—anecdotally including Yemen—that it has generated a phenomenon known in English as "beeping" (Donner, 2007), in which a person with little or no airtime on their account to make a call can dial, ring, and then hang up before the call connects, fully expecting the missed call will be interpreted as a message saying "I would like to speak with you but I do not have airtime, so please call back." Additionally, and less intuitively, wealth is negatively associated across varied societies with clustering coefficient: poorer people tend to have calling networks that are closer to being a clique (complete graph). We also use estimates of kilometer grid-square economic productivity from Ghosh et al. (2010) (Accessed via Goodman et al., 2019), which combines VIIRs nighttime light data and LandScan population density estimates to measure economic output per capita. We link tower locations to the corresponding grid-square in the Goodman et al. (2019) data.

E.3 Mobility

We calculate mobility metrics from CDRs using the `bandicoot` library (de Montjoye et al., 2016). We calculate mobility metrics first for the months of pre-occupation baseline data (January-March 2011), and then for each week of data during the conflict (March 2011-June 2012). The key measure is radius of gyration (RoG, a measure of how wide an individual's "typical" range of movement is over a period of time). We separately measure RoG to estimate an individual's baseline, pre-conflict behavior, and again as a weekly statistic that mobility behavior during conflict. All mobility data is aggregated to the level of cell towers, which is more likely to capture coarse movements than quotidian movements within a home location. Subscribers can theoretically travel up to tens of miles (though usually less) without changing towers, and conversely, subscribers in an area covered by two towers might switch towers without moving. RoG, unlike other mobility measures, accounts for the second quirk.

Some CDR studies like Gonzalez et al. (2008) note an inverse relationship in the short term between mobility and cell activity, which makes it challenging to use CDRs for fine-grained location track information. Intuitively this relationship makes sense: people communicate more at rest than on the move, making their communication locations an incomplete record of the locations they may have visited between CDRs. Our data structure avoids some of this problem, though, by aggregating records to the week level, and calculating central tendencies and statistics at that level. We further address the possible confounding caused by the relationship between cellphone activity and mobility by including prior mobility as a covariate in our models.

E.4 Religious Practice

We use proxies for subscriber religiosity (more precisely, adherence to the Islamic practice of daily prayer at specified times) as a final set of possible predictors. Civilians' religious practice is potentially an important alternative explanation for migration in the Yemeni context because one of the major parties to the conflict we study, AQAP, defines itself as a religious fundamentalist movement. A potential counter-argument to our claims about social network centrality, therefore, could be that migration is actually driven by an individual's ideological distance from the armed group occupying territory where they live. Since AQAP was notably lax in its enforcement of social practices and religious laws consistent with the group's ideology in the early 2010s in Yemen (Barfi, 2010), we do not expect that low rates of daily prayer would constitute a major push factor in this case. All the same, recent research on migration out of territory held by the Islamic State, suggests that approval of ISIS governing practices is associated at the 10% significance level with migration behavior(Revkin, 2019, Table 1). Regularity of religious practice is hardly a perfect proxy for approval of Salafism—the vast majority of observant Muslims around the world are not ideologically aligned with AQAP, but absent survey measures of attitudes, measuring regularity of prayer is one way to identify evidence that an individuals' pattern of religous practice would be *inconsistent* with the stated ideology of an occupying armed group. We use a measure of religiosity developed by Bozcaga et al. (2019) in the CDR context and also used by Dube et al. (2022) and Blaydes and Linzer (2012) , which focuses on calling behavior within specific temporal windows. We estimate our main models with "religiosity" measures composed from calling behavior in three distinct sets of windows, proposed by the above-cited authors. First, the models in the main text measure religiosity by comparing calling behavior on Fridays in the hour before mid-day prayer, and during the days of religious holidays. We calculate the number and duration of calls made by each subscriber in the hour before Friday mid-day prayers. Scaling the count and duration metrics provides a proxy for religiosity vis a vis other callers—those who are calling more in these key times are, per Bozcaga et al. (2019), relatively more religious. Second, we use the same procedure but apply it to the hour-long window before daily Maghrib prayer (evening) as it shifts seasonally with the time of sunset. Finally, we re-calculate the same measure using calling behavior in the hour before Fajr prayer (early morning). Calls in the hour before Fajr (ostensibly to coordinate prayer) are simultaneously the highest bar for religious adherence and the most discriminating measure. Even in the winter when sunrise is latest, Fajr prayers in Yemen commence no later than about 5:00am, and, consequently, are the least-observed of the five daily prayers Results from the additional measures are shown below in Appendix I. Our results are not sensitive to the particular prayer that is used to measure religious adherence.

E.5 Networks

The network statistics we describe in the body of the paper (Section 5.2) are calculated based on a graph composed of an individual subscribers' calls and texts in the period immediately before the March 2011 AQAP occupation of Abyan. We construct network graphs for every subscriber who we define as a resident of Abyan Governorate based on their continuous presence in the governorate for a period of four weeks in January 2011.[22] The Call Detail Records we use to create the individual graphs include *any* call or text where the Abyan resident is involved. As such, many Abyan residents' networks include alters who are not themselves residents of Abyan. Alters who are/not residents of Abyan are not distinguished when calculating network centrality statistics, but all models in Section 6 include covariates for mean call distance, such that network centrality results can be interpreted as the marginal association between social centrality and migration behavior, after adjusting for differences in the geographic spread or compactness of any individual's network of social alters.

We also briefly relate our measures to others used in the political science literature. Betweenness centrality is one common conceptual notion of centrality. Padgett and Ansell (1993), in a seminal

[22] We end the residency determination period 6 weeks before the start of the occupation to avoid missing individuals who might pre-emptively migrate as the conflict is starting but before their area of residence experiences violence.

historical study of the networks of Renaissance Florentine families, show that social influence depends heavily on betweenness centrality: how often a particlar node appears on the shortest path connecting two randomly selected nodes. Nodes with high betweenness may or may not have many connections, but they are critical for making connections between others. They have high "bridging" potential in other words (Berman, 1997). This is related to the intuition behind PageRank, one of our main measures, which itself is a directed graph variant of another commonly used metric: eigenvector centrality, which rates the "centrality" of a node based on the centrality of its neighbors. In plain terms, eigencentrality posits that an important element of relative social centrality is not difference in the length of contact lists, but difference in the quality of contact lists. Cruz et al. (2017) show that eigencentrality is associated with political importance or membership in important families in a study of the Philippines.

F Summary Statistics

Variable	Complete	Mean	SD	Min.	25%	Med.	75%	Max
Call Duration (preocc, avg)	0.959	166.095	178.714	2.000	58.769	99.388	200.806	1883.000
# Contacts (preocc, avg)	0.959	5.660	6.060	1.000	2.000	3.670	6.890	102.000
% Initiated (preocc, avg)	0.960	0.079	0.151	0.000	0.000	0.000	0.111	1.000
Call Distance km. (preocc, avg)	0.957	57.600	69.200	0.000	15.600	34.7000	73.100	908.000
Friday Prayer Calls (preocc, avg)	1.000	0.606	2.473	0.000	0.000	0.000	0.000	74.000
Radius of Gyr. (preocc, avg)	0.959	4.254	9.084	0.000	0.000	1.439	4.854	459.938
Pagerank (preocc)	0.960	1.32e-5	1.55e-5	9.33e-7	4.50e-6	8.90e-6	1.68e-5	0.002
Clust. Coef. (preocc)	0.960	0.093	0.141	0.000	0.013	0.051	0.111	1.000
Deg. Centrality (preocc)	0.960	9.48e-5	1.11e-4	4.34e-6	2.60e-5	6.07e-5	1.21e-4	0.004
Latitude (tower, lag)	0.889	13.554	0.663	12.685	13.134	13.359	13.883	17.796
Longitude (tower, lag)	0.889	45.279	0.851	42.675	44.995	45.298	45.382	53.095
Log GDP (tower, lag)	0.888	-0.822	2.056	-7.428	-1.878	-1.833	0.286	4.538
Pre-2010 GED Events (tower, lag)	0.888	73.421	105.303	0.000	0.000	16.000	214.000	408.000
GED Event Count (tower, lag)	0.979	0.109	0.537	0.000	0.000	0.000	0.000	11.000
# Records (person, lag)	0.961	24.835	43.605	0.000	1.000	9.000	29.000	1300.000
Radius of Gyr. (person, lag)	0.736	6.442	17.051	0.000	0.000	1.178	5.068	1179.088
% Pareto Int. (person, lag)	0.752	0.220	0.154	0.000	0.103	0.188	0.333	0.800

Table A.3: Summary statistics (and missingness) for predictor variables. Variables are de-meaned and scaled in model estimations.

Variable	Complete	Mean	Logical Vals.
In Abyan?	1.000	0.546	T: 1,568,078 F: 1,302,180

Table A.4: Dependent variable summary for time series estimations.

Variable	Complete	Min.	Max.	Med.	Unique Val.
Date	1.000	2011-01-30	2012-06-24	2011-10-16	74

Table A.5: Date summaries for time series estimation.

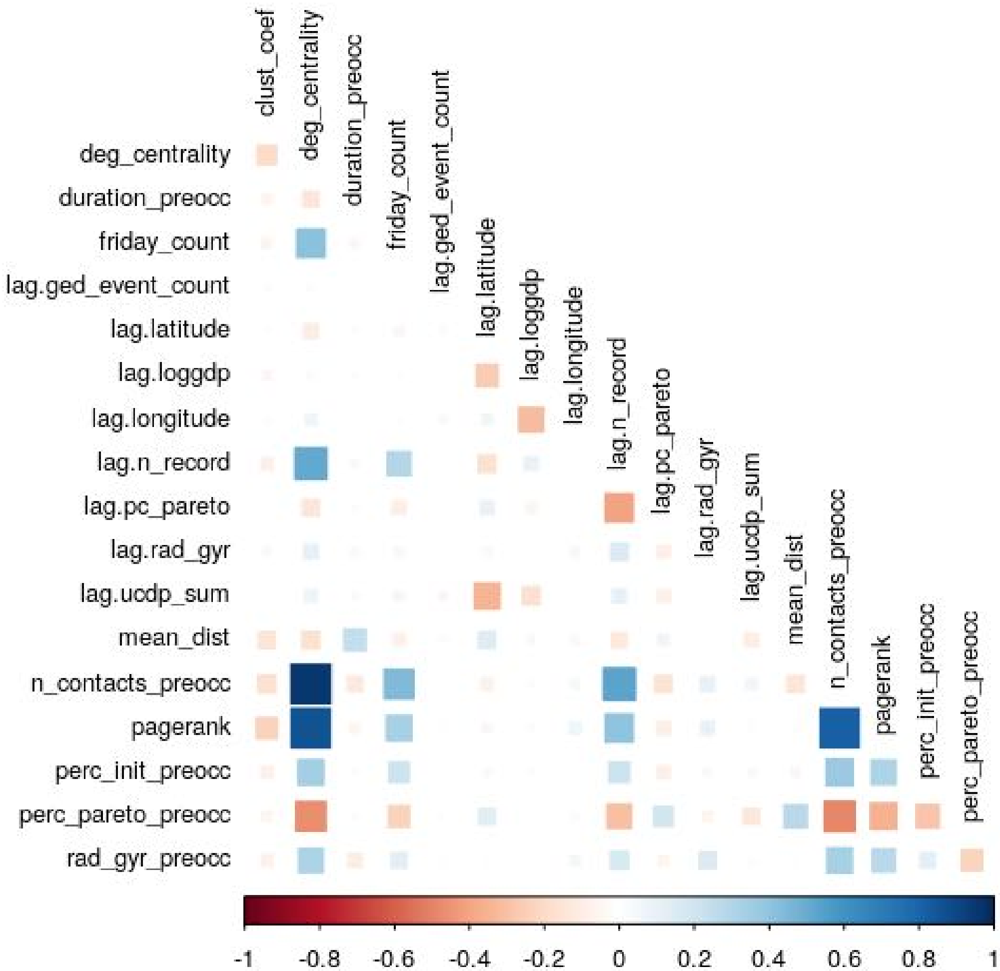

Figure A.4: Correlation of predictor variables (raw, un-transformed values). The highest bivariate correlation is 0.93.

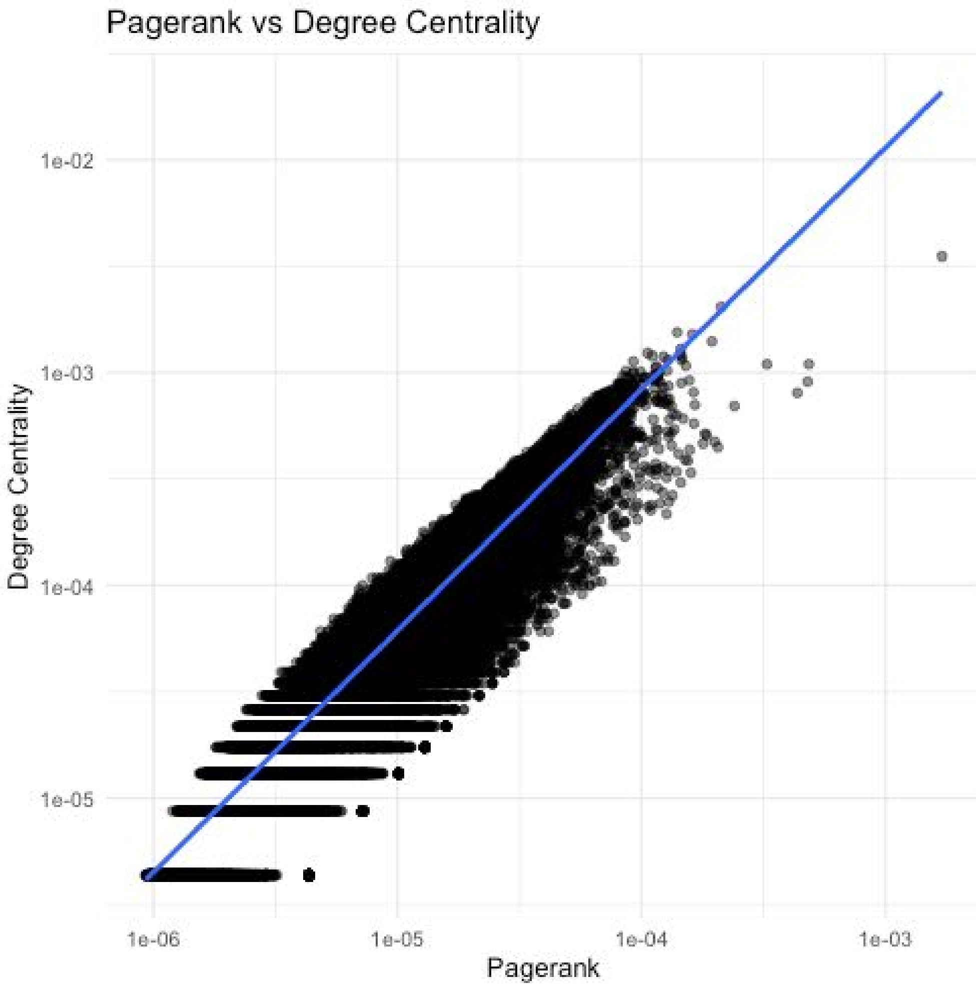

Figure A.5: Correlation of PageRank and Degree Centrality, two network measures which we juxtapose in the hypotheses and results. The measures are somewhat highly correlated $r = .83$, which we expect because they measure conceptually related quantities. The building blocks of Degree Centrality, in fact, are one component of PageRank. PageRank generally measures influence by the quantity of a node's out-directed edges (a main component of degree centrality), plus the quantity of the out-directed edges of all the original node's connections (a component of those first-degree connections' degree centrality). We expect that the two metrics are therefore correlated but not collinear. When we estimate models using both PageRank and Degree Centrality, we use the phrase "Influence" to describe the PageRank coefficient, because a node's influence (i.e. the embeddedness of that node's *neighbors*) is what the marginal association with PageRank measures after holding degree centrality constant.

G Subscriber Anonymity, and Data Ethics

The CDR data we analyze in this paper come from one of Yemen's three major cellphone providers circa 2010. As noted elsewhere in the SI, each of the three major providers was well-established and nationally prevalent in the time period represented by our data. Data were shared directly with [Author 2] by the company. The initial agreement between [Author 2] and the telecom company allows that the data may be used for social scientific inquiry at the researcher's discretion, but that the raw data may not be publicly shared, and that the providing company may not be named in reports or publications.[23] The data sharing agreement does not allow for the telecom company to weigh in on the scope, analysis, or findings from research undertaken using the data, and does not require that results are reported back to the telecom company—given political upheaval in Yemen between the end of the data and the present day, the results of our research are likely of little present-day practical value to corporate or government interests in Yemen.

Our data, to be clear, are "anonymous" in that no personally identifiable information is represented in the data: phone numbers were replaced with random digit unique identifiers *by the cellphone company, prior to researcher access*. The research team has never had access to any personally identifiable information from the subscribers represented in the dataset. As we describe below, this does not obviate concerns about anonymity and human subject protections—our study was approved as human subjects research after review by the IRB at [home institution]. We believe, however, that three features of our data plus additional steps we take in analysis and presentation, make privacy protections quite robust. First, anonymity at the source separates our data from other sources of remote-sensing or communication/internet behavior data that social scientists are increasingly using to study conflict (Abrahams and Greenwald, 2021). Second, through a combination of data limitations inherent in 2G cellphone records and additional privacy preserving choices in our analysis, we do not ask or answer questions about hyper-local variation in migration that might edge toward probabilistic identification of individual households. Third, our data are now over a decade old. Recent oddities in the telecom market in Yemen have led to phone companies re-assigning or duplicating

[23] This happens to be a benefit for privacy preservation!

dormant numbers without the knowledge or consent of the original users (see: globalvoices.org). Even if it were possible to identify particular cell phone numbers in our data, it is not clear that the number could be subsequently linked to an individual user.

Of course, anonymous data are not guaranteed to stay that way. A number of large datasets that are supposed to be anonymous have been "de-anonymized" in recent years, including datasets built from CDRs. de Montjoye et al. (2013) show that as few as four call detail records with time and location stamps (cell tower locations)[24] can uniquely identify a caller in anonymous data, such that an individual location trace can be matched to a specific person using contextual knowledge, and that increasing the spatial and temporal coarseness of the data can decrease the threat of re-identification only if the resolution of the data is lowered substantially. We follow the recommendations implied by de Montjoye et al. (2013) and also note some differences between their data and ours which should further complicate any efforts to de-anonymize the CDRs.

The "four points" finding comes from cell data in a Western European country, and is based on antenna density that is substantially higher than we see in Abyan. Tower density is important for anonymity: if the CDR dataset can distinguish between someone making a call from 6th avenue and someone making a call from 5th avenue, de-anonymization is easier than if the call is only known to be coming from "one of the middle avenues" in Manhattan. Tower density in Abyan is very low compared to highly urbanized Western European countries at the same time, meaning that even individual tower IDs provide less information about the location of a caller.

Though re-identification is harder in our data simply due to the nature of the communication infrastructure in Abyan, we take further steps to protect anonymity. All the data products we use are aggregations of multiple CDRs to a level of coarseness which far exceeds the coarsest level tested in de Montjoye et al. (2013). They investigate lowering temporal resolution up to 15 hours, our smallest

[24] To be clear, all our location measures in the paper are proxies based on the location of the cell tower through which a user's call connects. Our location measures are therefore orders of magnitude less precise than the sort of GPS trace data from smartphones that has raised privacy concerns in recent years (Thompson and Warzel, 2019).

unit of analysis is a whole week (11.2 times coarser). They investigate coarsening tower resolution up to the seven-tower level, we do not report each tower that a subscriber visited, but only the ”most frequent” over the course of a seven day period. Further, none of the analyses in our paper rely on extracting or isolating individuals or small groups of subscribers. We will not release the underlying CDR data, a greater de-anonymization risk, as it is not necessary to replicate our analyses.

A corresponding source of non-anonymized data is another key point of leverage for de-anonymization that is relatively absent in the case we study. Re-identifying anonymous communication, mobility, and social network data typically requires one of two points of leverage. The first, as in the troubling de-anonymization of data from Lewis et al. (2008), is contextual knowledge. In the case of our data, the necessary contextual knowledge would be quite arcane even if we presented profiles of individual subscribers in the data because our measures of location (cell tower) do not resolve to home addresses or even necessarily to neighborhoods and even our simplest measures of communication behavior are probably difficult to recognize “in the wild.”[25] The second possible point of leverage for de-anonymization would be a second, personally identifiable network dataset that could be compared to our network graphs. Information from Facebook, for instance, could theoretically be compared to our data to make guesses about the identities of individuals in our data based on their communication and social ties. Narayanan and Shmatikov (2009) demonstrate that this type of de-anonymization is a particular threat in sparse data, which include some social networks. Fortunately for the purposes of our study, suitable second sources of network data are extremely hard to find for circa 2010 residents of Abyan. As we note in the paper, internet penetration in Yemen generally was extremely low in 2010, so additional datasets characterizing the communication and mobility behavior of Abyan residents are very unlikely to exist. Without additional sources of network data to compare to the anonymized cell phone data we analyze the prospect of de-anonymizing data in this way is very slim.

[25]The authors, for instance, cannot report without checking their phones how many calls they made/received in the past week, and what proportion they initiated.

H Geographic Measurement Proxies

Location data in our paper are measured using the cell tower through which a call or text connected as a proxy for subscriber location. This is standard in the CDR literature, but it is still worth discussing how it works, and what its drawbacks are. The technical limits on how far a cell phone can be from a tower it is connecting through are surprisingly high: the older GSM system used in Yemen at the time our data were collected has a technical limit of 35km range. The functional range is almost always smaller, but varies widely based on a) topography, b) the built environment, and c) intentional "detuning" by cellular providers in order to prevent tower-to-tower interference while sharing a relatively narrow frequency range.

Range is only one part of the equation, however. Cell phones often, but not always, connect through the tower that is physically closest. When multiple towers are available, networks route phones to one specific tower with input from the handset's signal strength measurements. Most users of cellphones know from personal experience that quality of connection is not smooth with respect to distance from some fixed point. Cell towers as a proxy for location, then, can be thought of as providing a hard *upper bound* for a subscriber's location—they cannot be situated more than 35km from the tower when they connect to it—but practically, people are usually much closer.

Finally, it is worth noting that there is a way to identify much more precise location using only cellular connection data (i.e. not GPS pings used in Pappalardo et al. (2023)), which is sometimes used by emergency services in countries including the United States to identify the location of callers in distress. This method relies on two pieces of data that are not recorded in our CDRs and cannot be reverse-engineered from it: 1) it requires the ID/location of a second tower that has exchanged a signal handshake with the phone (is in range of the phone), and 2) it requires the signal angle of arrival to both towers. These data support the method colloquially called "triangulation," which we cannot use.

I Additional Model Figures/Tables

I.1 Model Specifications

Linear Random Effects model (LMER) specification for results in Figure 3a:

$$Y_i = \alpha + \xi K + \delta J + \beta \mathbf{X}_i + \varepsilon_i$$

Where:

- α = Global intercept
- Y_i = Proportion of time outside Abyan
- K = Tower dummies
- J = District dummies
- X includes: Number of contacts, call duration, % initiated interactions, % pareto interactions, mean call distance, Friday prayer calls, radius of gyration, Pagerank, clustering coefficient, degree centrality, Latitude, Longitude, Lat/Lon polynomial, pre-conflict violence events, tower GDP.

Logit Random Effects model specification for results in Figure 3b:

$$logit(P(Y_{it} = 1)) = \alpha_i + \gamma_t + \beta \mathbf{X}_{it}$$

Where:

- Y_{it} = Binary indicator of weekly presence in Abyan
- i = Subscriber
- t = Time period
- $logit(P(Y_{it} = 1)) = ln(\frac{P(Y_{it}=1)}{1-P(Y_{it}=1)})$, the log odds transformation
- γ_t = Random effect for time period, assumed to be normally distributed $\gamma \sim N(0, \sigma_t^2)$
- α_t = Random effect for subscriber, assumed to be normally distributed $\alpha \sim N(0, \sigma_i^2)$
- X includes: Number of contacts, call duration, % initiated interactions, % pareto interactions, mean call distance, Friday prayer calls, radius of gyration, Pagerank, clustering coefficient, degree centrality, Latitude, Longitude, Lat/Lon polynomial, pre-conflict violence events, tower GDP.

I.2 Figures and Tables

Parameter	Estimate	Std. Err.	P-Value	2.5%	97.5%
Intercept	4.5149942	29.3463347	0.8777264	-53.0027648	62.0327533
# Contacts (preocc, avg)	-0.0671430	0.0063389	0.0000000	-0.0795670	-0.0547190
Call Duration (preocc, avg)	-0.0036961	0.0025880	0.1532422	-0.0087684	0.0013763
% Initiated (preocc, avg)	0.0094660	0.0150928	0.5305380	-0.0201154	0.0390474
% Pareto Int. (preocc, avg)	-0.2324781	0.0142842	0.0000000	-0.2604747	-0.2044816
Mean Call Dist (preocc)	-0.0028714	0.0031991	0.3694176	-0.0091415	0.0033987
Friday Prayer Calls (preocc)	0.0029521	0.0019618	0.1323715	-0.0008929	0.0067971
Radius of Gyr. (preocc, avg)	-0.0051375	0.0024819	0.0384548	-0.0100019	-0.0002730
Pagerank (preocc)	-0.0098330	0.0047560	0.0386858	-0.0191545	-0.0005115
Clustering Coef. (preocc)	0.0092493	0.0032083	0.0039402	0.0029611	0.0155374
Deg. Centrality (preocc)	0.1086211	0.0072390	0.0000000	0.0944329	0.1228092
Latitude	-0.3725637	2.0936730	0.8587643	-4.4760874	3.7309600
Longitude	-0.1176713	0.6578136	0.8580300	-1.4069622	1.1716196
Lat/Lon Poly (tower)	0.0105029	0.0469452	0.8229703	-0.0815081	0.1025139
Pre-2010 UCDP Events (tower)	0.0560421	0.0030637	0.0000000	0.0500374	0.0620468
Log GDP (tower)	-0.0163719	0.0205024	0.4245599	-0.0565559	0.0238121

Table A.6: Raw coefficient estimates corresponding to Figure 3a. Signs are reversed in the plot in order to show marginal associations with time outside of Abyan, not Inside (the dependent variable, shown in Table A.4 takes a value of 1 for remaining.

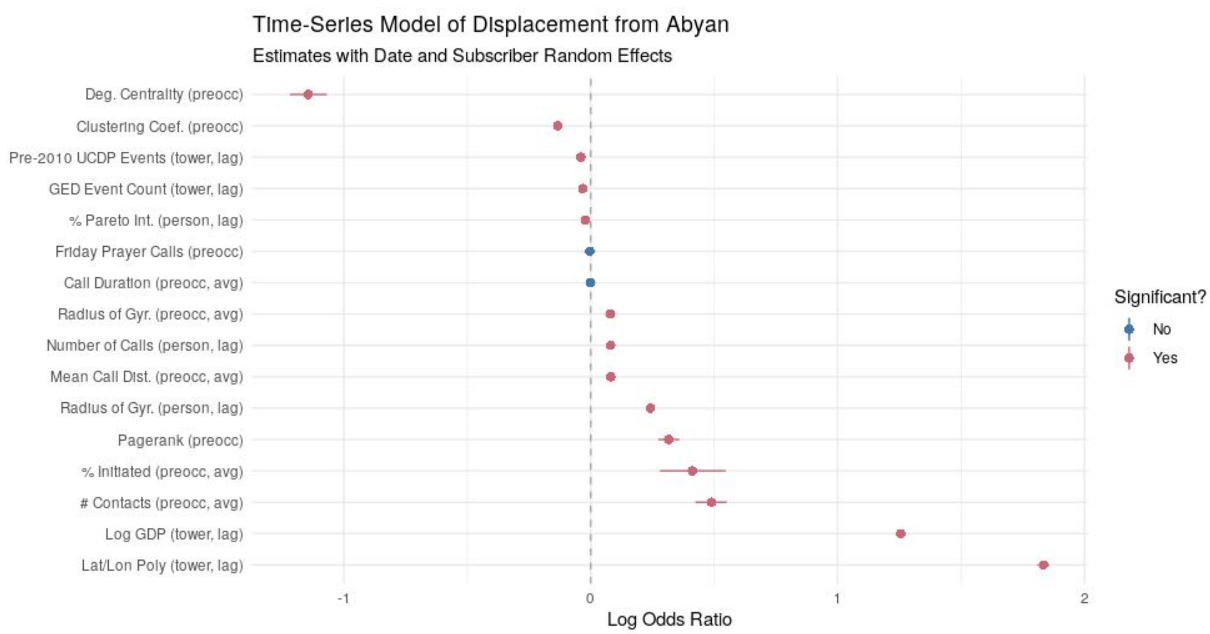

Figure A.6: Raw log odds ratio parameter estimates from the HMC Estimation presented in Figure 3b. Log odds ratios here are associated with a one standard deviation change in the value of the predictor, as all numerical predictors (save latitude and longitude) are demeaned and scaled to facilitate better exploration of the posterior density. Error bars show 99% credible intervals, calculated from stan's sampling of the posterior density function.

Parameter	mean	mcse	sd	10%	50%	90%	n_eff	Rhat
(Intercept)	-1213.265	0.086	8.092	-1223.777	-1213.175	-1202.879	8764	1.000
Number of contacts (Preocc, avg)	-0.489	0.001	0.032	-0.530	-0.489	-0.447	1179	1.007
Call duration (Preocc, avg)	0.001	0.000	0.010	-0.012	0.001	0.013	1350	1.004
% Initiated Conv. (Preocc, avg)	-0.412	0.002	0.067	-0.499	-0.412	-0.328	1405	1.005
Average call distance (Preocc)	-0.081	0.000	0.010	-0.094	-0.081	-0.069	1736	1.002
Friday prayer calls (Preocc)	0.003	0.000	0.011	-0.011	0.003	0.018	1050	1.006
Radius of gyration (Preocc, avg)	-0.080	0.000	0.010	-0.093	-0.080	-0.066	1513	1.002
Pagerank (Preocc)	-0.316	0.001	0.022	-0.344	-0.316	-0.288	1341	1.006
Clust. Coef (Preocc)	0.134	0.000	0.009	0.122	0.134	0.146	1812	1.002
Degree Centrality (Preocc)	1.144	0.001	0.038	1.094	1.144	1.193	967	1.009
Latitude (tower,lag)	80.764	0.006	0.555	80.053	80.757	81.483	8855	1.000
Longitude (tower,lag)	27.476	0.002	0.181	27.244	27.473	27.711	8817	1.000
Lat*Lon Polynomial (tower,lag)	-1.834	0.000	0.012	-1.850	-1.833	-1.818	8856	1.000
Log GDP (tower,lag)	-1.256	0.000	0.005	-1.263	-1.256	-1.249	9813	1.000
Pre-2010 GED Events (tower, lag)	0.040	0.000	0.005	0.034	0.040	0.046	8977	1.000
GED Event count (tower, lag)	0.032	0.000	0.003	0.029	0.032	0.035	9854	0.999
Number of calls (lag)	-0.081	0.000	0.004	-0.086	-0.081	-0.076	9086	1.000
Radius of gyration (lag)	-0.242	0.000	0.003	-0.246	-0.242	-0.238	9433	1.000
Percent Pareto Interactions (lag)	0.021	0.000	0.003	0.017	0.021	0.025	10041	1.000

Table A.7: Raw parameter estimates and model diagnostics from Hamiltonian Monte-Carlo (HMC) estimation of a multilevel logistic regression with time- and unit- varying intercepts shown in Figure 3b. Key diagnostics for exploration of the posterior—n_eff and Rhat—do not indicate issues with the model's exploration of the density function. Gelman et al. (2013) and The Stan Development Team (2021) suggest not to interpret models with Rhat values > 1.05 or n_eff values below $5m$, where m is the number of chains after splitting.

Parameter	Median APE	0.5 Percentile Est.	99.5 Percentile Est.
Number of contacts (Preocc, avg)	-0.0691483	-0.0843806	-0.0547887
Call duration (Preocc, avg)	0.0000848	-0.0023765	0.0025607
% Initiated Conv. (Preocc, avg)	-0.0073333	-0.0102367	-0.0048960
Average call distance (Preocc)	-0.0101628	-0.0133024	-0.0073659
Friday prayer calls (Preocc)	0.0000000	0.0000000	0.0000000
Radius of gyration (Preocc, avg)	-0.0071395	-0.0094888	-0.0051076
Pagerank (Preocc)	-0.0436540	-0.0534913	-0.0346964
Clust. Coef (Preocc)	0.0135088	0.0107224	0.0166528
Degree Centrality (Preocc)	0.1529770	0.1253409	0.1821949
Latitude (tower,lag)	0.9999997	0.9999996	0.9999998
Longitude (tower,lag)	0.9753132	0.9681325	0.9807699
Lat*Lon Polynomial (tower,lag)	-1.0000000	-1.0000000	-1.0000000
Log GDP (tower,lag)	-0.4599450	-0.4733246	-0.4338020
Pre-2010 GED Events (tower, lag)	0.0126933	0.0093694	0.0164930
GED Event count (tower, lag)	0.0000000	0.0000000	0.0000000
Number of calls (lag)	-0.0101730	-0.0122323	-0.0083014
Radius of gyration (lag)	-0.0119649	-0.0139862	-0.0099782
Percent Pareto Interactions (lag)	0.0043684	0.0028920	0.0060301

Table A.8: Average partial effect estimates corresponding to Figure 3b. Signs are reversed in the plot in order to show average partial effect on probability of leaving, not remaining (the dependent variable, shown in Table A.4 takes a value of 1 for remaining.

Parameter	Median	2.5%	97.5%
Intercept	-1213.1747960	-1229.0921842	-1197.4177928
# Contacts (preocc, avg)	-0.4889985	-0.5511139	-0.4241460
Call Duration (preocc, avg)	0.0006676	-0.0188245	0.0201217
% Initiated (preocc, avg)	-0.4116512	-0.5463700	-0.2807507
Mean Call Dist. (preocc, avg)	-0.0813945	-0.1013038	-0.0615684
Friday Prayer Calls (preocc)	0.0031660	-0.0181898	0.0247066
Radius of Gyr. (preocc, avg)	-0.0799217	-0.1000910	-0.0592208
Pagerank (preocc)	-0.3156982	-0.3594177	-0.2741760
Clustering Coef. (preocc)	0.1337839	0.1156159	0.1521835
Deg. Centrality (preocc)	1.1440058	1.0690886	1.2182049
Latitude (tower, lag)	80.7570660	79.6773096	81.8541257
Longitude (tower, lag)	27.4731782	27.1199930	27.8297528
Lat/Lon Poly (tower, lag)	-1.8334745	-1.8580053	-1.8092751
Log GDP (tower, lag)	-1.2562643	-1.2669083	-1.2456920
Pre-2010 UCDP Events (tower, lag)	0.0401698	0.0313599	0.0490349
GED Event Count (tower, lag)	0.0320105	0.0270127	0.0372811
Number of Calls (person, lag)	-0.0810619	-0.0889503	-0.0731589
Radius of Gyr. (person, lag)	-0.2418761	-0.2475830	-0.2362507
% Pareto Int. (person, lag)	0.0210811	0.0144517	0.0276498

Table A.9: 95% credible intervals corresponding to figure A.6, and the parameter estimates/diagnostics in Table A.7. Signs are reversed in the plot in order to show effects on probability of leaving, not remaining (the dependent variable, shown in Table A.4 takes a value of 1 for remaining).

	Dependent variable:	
	Weeks Inside Abyan (proportion)	
	(1)	(2)
# Contacts (preocc, avg)	-0.145^{***}	0.039^{***}
	(0.003)	(0.001)
Call Duration (preocc, avg)	-0.0001	-0.004^{***}
	(0.001)	(0.001)
% Initiated (preocc, avg)	-0.030^{***}	-0.071^{***}
	(0.007)	(0.008)
% Pareto Int. (preocc, avg)	-0.201^{***}	-0.206^{***}
	(0.006)	(0.006)
Mean Call Dist (preocc)	-0.003^{**}	-0.011^{***}
	(0.001)	(0.001)
Friday Prayer Calls (preocc)	0.004^{***}	0.006^{***}
	(0.001)	(0.001)
Radius of Gyr. (preocc, avg)	-0.008^{***}	-0.013^{***}
	(0.001)	(0.001)
Latitude	0.283	-0.095
	(0.409)	(0.428)
Longitude	0.071	-0.039
	(0.135)	(0.141)
Lat/Lon Poly (tower)	-0.005	0.003
	(0.009)	(0.010)
Pre-conflict UCDP Events (tower)	0.029	0.038
	(0.024)	(0.027)
Log GDP (tower)	-0.021	-0.018
	(0.013)	(0.014)
Pagerank (preocc)	-0.040^{***}	
	(0.002)	
Clustering Coef. (preocc)	0.007^{***}	
	(0.001)	
Deg. Centrality (preocc)	0.230^{***}	
	(0.004)	
Constant	-3.261	1.897
	(6.080)	(6.355)
Observations	51,835	51,835
Log Likelihood	$-1,399.747$	$-3,351.148$
Akaike Inf. Crit.	2,837.493	6,734.297
Bayesian Inf. Crit.	3,005.754	6,875.990
Marginal R^2	0.130224	0.07309827
Note:	*p<0.1; **p<0.05; ***p<0.01	

Table A.10: LMER results comparing models with and without centrality statistics. Note the 1.781493x increase in Marginal R^2 (Model 1 over Model 2) for the inclusion of the network stats. Including versus omitting centrality measures also meaningfully changes some other associations like the number of contacts, and call duration.

	Dependent variable:			
	Weeks Inside Abyan (proportion)			
	(1)	(2)	(3)	(4)
# Contacts (preocc, avg)	−0.136***	0.013***	−0.145***	−0.145***
	(0.003)	(0.002)	(0.003)	(0.003)
Call Duration (preocc, avg)	−0.002	−0.004***	−0.001	−0.0001
	(0.001)	(0.001)	(0.001)	(0.001)
% Initiated (preocc, avg)	−0.037***	−0.070***	−0.031***	−0.030***
	(0.007)	(0.008)	(0.007)	(0.007)
% Pareto Int. (preocc, avg)	−0.209***	−0.209***	−0.206***	−0.201***
	(0.006)	(0.006)	(0.006)	(0.006)
Mean Call Dist (preocc)	−0.008***	−0.013***	−0.004***	−0.003**
	(0.001)	(0.001)	(0.001)	(0.001)
Friday Prayer Calls (preocc)	0.005***	0.006***	0.004***	0.004***
	(0.001)	(0.001)	(0.001)	(0.001)
Radius of Gyr. (preocc, avg)	−0.008***	−0.012***	−0.009***	−0.008***
	(0.001)	(0.001)	(0.001)	(0.001)
Latitude	0.252	−0.055	0.286	0.283
	(0.412)	(0.426)	(0.410)	(0.409)
Longitude	0.063	−0.028	0.073	0.071
	(0.136)	(0.141)	(0.135)	(0.135)
Lat/Lon Poly (tower)	−0.005	0.002	−0.005	−0.005
	(0.009)	(0.009)	(0.009)	(0.009)
Pre-treatment UCDP Events (tower)	0.030	0.038	0.029	0.029
	(0.025)	(0.027)	(0.025)	(0.024)
Log GDP (tower)	−0.021	−0.018	−0.021	−0.021
	(0.013)	(0.014)	(0.013)	(0.013)
Deg. Centrality (preocc)	0.185***		0.229***	0.230***
	(0.003)		(0.004)	(0.004)
Clustering Coef. (preocc)				0.007***
				(0.001)
Pagerank (preocc)		0.033***	−0.040***	−0.040***
		(0.002)	(0.002)	(0.002)
Constant	−2.847	1.375	−3.320	−3.261
	(6.114)	(6.333)	(6.085)	(6.080)
Observations	51,835	51,835	51,835	51,835
Log Likelihood	−1,598.931	−3,176.096	−1,414.018	−1,399.747
Akaike Inf. Crit.	3,231.863	6,386.193	2,864.037	2,837.493
Bayesian Inf. Crit.	3,382.412	6,536.742	3,023.442	3,005.754
Note:				*p<0.1; **p<0.05; ***p<0.01

Table A.11: These linear model results in columns 1-3 show the effects of leaving out individual centrality measures. As expected, degree centrality and pagerank are somewhat highly correlated ($cor = 0.83$) in the data because they measure related concepts. Comparing the models, we can interpret the models as showing that the part of Pagerank that is related to degree centrality (first degree network breadth) is a reasonably large part of Pagerank's total effect, and is associated with *remaining*. The elements of Pagerank that are **not** captured in degree centrality—the strength of the networks of a node's contacts—are substantively meaningful even after controlling for degree centrality, and predict leaving. This motivates our decision in the main text to interpret page rank as "influence" and degree centrality as "breadth." Column 4 shows the same results as the fully saturated model in Table A.10 .

I.3 Additional Religious Practice Proxies

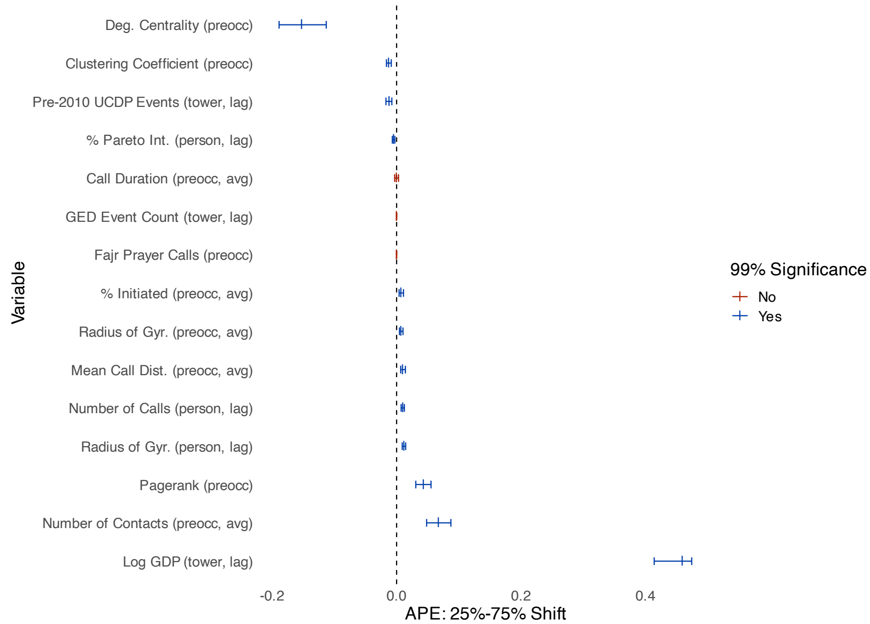

Figure A.7: Average partial effect estimates from a `stan` model re-estimating the model in Figure 3b, but substituting the Bozcaga et al. (2019) measure of religiosity for a measure based on calls before daily Fajr prayer, performed between dawn and sunrise (See more detail in Appendix E). Results and their substantive interpretation as described in Section 6 hold with this alternative measure of religiosity. Tabular results corresponding to this figure are in Table A.12

Table A.12: APE estimates corresponding to Figure A.7. Signs are reversed in the plot in order to show changes in the probability of leaving, not remaining.

Parameter	Median Partial Effect	0.5%	99.5%
Number of Contacts (preocc, avg)	-0.0674	-0.0873	-0.0483
Call Duration (preocc, avg)	0.0000	-0.0031	0.0031
% Initiated (preocc, avg)	-0.0071	-0.0111	-0.0040
Mean Call Dist. (preocc, avg)	-0.0100	-0.0142	-0.0066
Fajr Prayer Calls (preocc)	0.0000	0.0000	0.0000
Radius of Gyr. (preocc, avg)	-0.0071	-0.0103	-0.0046
Pagerank (preocc)	-0.0433	-0.0553	-0.0308
Clustering Coefficient (preocc)	0.0126	0.0085	0.0163
Deg. Centrality (preocc)	0.1523	0.1130	0.1888
Log GDP (tower, lag)	-0.4593	-0.4742	-0.4140
Pre-2010 UCDP Events (tower, lag)	0.0119	0.0076	0.0171
GED Event Count (tower, lag)	0.0000	0.0000	0.0000
Number of Calls (person, lag)	-0.0101	-0.0125	-0.0073
Radius of Gyr. (person, lag)	-0.0121	-0.0146	-0.0091
% Pareto Int. (person, lag)	0.0045	0.0027	0.0068

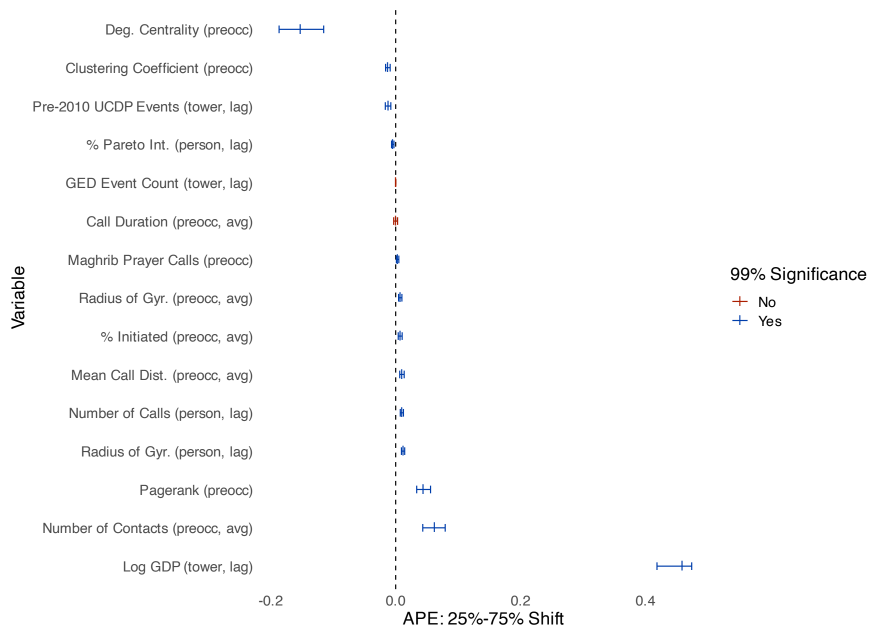

Figure A.8: Average partial effect estimates from a `stan` model re-estimating the model in Figure 3b, but substituting the Bozcaga et al. (2019) measure of religiosity for a measure based on calls before daily Maghrib prayer, performed directly after sunset (See more detail in Appendix E). Results and their substantive interpretation as described in Section 6 hold with this alternative measure of religiosity. Tabular results corresponding to this figure are in Table A.13

Table A.13: APE estimates corresponding to Figure A.8. Signs are reversed in the plot in order to show changes in the probability of leaving, not remaining

Parameter	Median Partial Effect	0.5%	99.5%
Number of Contacts (preocc, avg)	-0.0624	-0.0793	-0.0434
Call Duration (preocc, avg)	-0.0001	-0.0031	0.0032
% Initiated (preocc, avg)	-0.0073	-0.0104	-0.0041
Mean Call Dist. (preocc, avg)	-0.0099	-0.0136	-0.0065
Maghrib Prayer Calls (preocc)	-0.0030	-0.0050	-0.0012
Radius of Gyr. (preocc, avg)	-0.0072	-0.0100	-0.0047
Pagerank (preocc)	-0.0442	-0.0558	-0.0336
Clustering Coefficient (preocc)	0.0126	0.0088	0.0162
Deg. Centrality (preocc)	0.1526	0.1153	0.1867
Log GDP (tower, lag)	-0.4588	-0.4740	-0.4184
Pre-2010 UCDP Events (tower, lag)	0.0120	0.0078	0.0166
GED Event Count (tower, lag)	0.0000	0.0000	0.0000
Number of Calls (person, lag)	-0.0100	-0.0124	-0.0074
Radius of Gyr. (person, lag)	-0.0120	-0.0144	-0.0092
% Pareto Int. (person, lag)	0.0046	0.0028	0.0066

I.4 Directed PageRank Results

As noted in Section 5.2, we use undirected PageRank in our main specifications because call direction in Yemen's prepaid market reflects airtime credit constraints that complicate the interpretation of edge direction as a measure of social importance Donner (2007). Here, we re-estimate the main model substituting *directed* PageRank. Directed PageRank scores each node by the PageRank of the nodes that point toward it (i.e., those that call or text it), weighting incoming connections by the importance of the sender (Page et al., 1998).

Table A.14 presents full posterior estimates. The directed PageRank coefficient is negative and precisely estimated (posterior mean $= -0.133$, 95% CI $[-0.166, -0.101]$, P($>$0) ≈ 0), consistent in sign and significance with the undirected PageRank coefficient in the main model (-0.31, 95% CI $[-0.36, -0.26]$), though smaller in magnitude (See Table A.7 for details). Degree centrality likewise remains strongly positive ($+0.970$, 95% CI $[0.901, 1.039]$). We interpret directed and undirected PageRank as capturing influence and reachability respectively; graph orientation does not alter the substantive conclusion about the associations of these two related social characteristics with migration propensity.

Parameter	Mean	MCSE	SD	10%	50%	90%	n_{eff}	$\hat{R}$
# Contacts (preocc, avg)	−0.483	0.002	0.030	−0.521	−0.483	−0.445	208	1.014
Call Duration (preocc, avg)	0.003	0.001	0.010	−0.010	0.003	0.015	357	1.004
% Initiated Conv. (preocc, avg)	−0.466	0.004	0.061	−0.543	−0.467	−0.386	225	1.019
Avg. Call Distance (preocc)	−0.104	0.000	0.010	−0.116	−0.104	−0.092	395	1.023
Fajr Prayer Calls (preocc)	−0.014	0.001	0.010	−0.027	−0.014	0.000	156	1.035
Directed PageRank (preocc)	−0.133	0.001	0.016	−0.154	−0.133	−0.112	187	1.034
Clustering Coef. (preocc)	0.133	0.001	0.009	0.122	0.133	0.145	276	1.006
Degree Centrality (preocc)	+0.970	0.003	0.035	0.925	0.970	1.014	167	1.021
Log GDP (tower, lag)	−1.254	0.000	0.005	−1.260	−1.254	−1.247	4303	1.001
UCDP Event Count (tower, lag)	0.039	0.000	0.004	0.034	0.039	0.045	4280	1.000
# Records (person, lag)	−0.079	0.000	0.004	−0.084	−0.079	−0.074	4709	1.001
Radius of Gyr. (person, lag)	−0.244	0.000	0.003	−0.248	−0.244	−0.240	9749	1.001
% Pareto Int. (person, lag)	0.022	0.000	0.003	0.018	0.022	0.026	7162	1.000
Latitude (tower, lag)	−1.506	0.000	0.009	−1.517	−1.506	−1.495	2571	1.001
Longitude (tower, lag)	2.238	0.000	0.012	2.223	2.238	2.253	1497	1.002

Table A.14: Posterior parameter estimates from the directed PageRank robustness model. Specification identical to the main time-series model except directed PageRank replaces undirected PageRank. All predictors demeaned and scaled. MCSE = Monte Carlo standard error; n_{eff} = bulk effective sample size; $\hat{R}$ = potential scale reduction factor.

I.5 Alternative Centrality Measures

Our primary measure of "influence" or "importance" in the social network of Abyan residents detailed in Figure 2 and Table 1 is the PageRank centrality measure (Page et al., 1998). As detailed above in Appendix E.5, PageRank is a widely used and relatively intuitive centrality measure for estimating the influence or importance of a particular node in a graph. One advantage of PageRank is its ability to capture indirect or recursive influence across multiple connections, which approximates how ideas or beliefs (like beliefs about a conflict environment) spread through a social world. Two of the major drawbacks of the PageRank algorithm are susceptibility to "sink nodes" or nodes with only in-directed edges, and instability in evolving or dynamic social networks (Ghoshal and Barabási, 2011).

The first issue of sink nodes is less plausible with a network built from CDR data than a network of web links. In the graph we analyze, nodes with an out-degree of zero actually have a slightly lower average PageRank Centrality score than nodes with an out degree greater than zero. They also have, on average, fewer in-directed edges than the graph as a whole. They are simply less connected, and do not seem to be skewing PageRank centrality in ways the literature identifies as problematic.

The second issue primarily pertains to evolving social networks, whereas we estimate centrality in a single time window to avoid confounding between displacement and other social behavior. As such, updated versions of the PageRank algorithm (PPR, HITS) are not likely to be well suited to our empirical application.

Because PageRank is only one way of capturing the underlying concept of "importance," though, we also evaluate the association between influential network position and probability of displacement from Abyan using alternate measures of social importance and influential network position. We use two measures—Harmonic centrality (Marchiori and Latora, 2000; Rochat, 2009) and Strength centrality (Barrat et al., 2004; Opsahl et al., 2010) that prioritize separate notions of "influence." Compared to PageRank, Harmonic centrality—a variant of closeness centrality—prioritizes reachability, which is a natural component of "importance" in networks that might share information about conflict risk. Strength centrality, in turn, measures cumulative edge weight, which captures direct influence as a function of how regularly two nodes connect. As the name suggests, this measure of centrality captures the average "strength" of a node's relationships, which is a separate aspect of being important in a network.

Results from models where PageRank is substituted for Strength or Harmonic centrality are presented in Figures A.10 and A.9 and Tables A.16 and A.15, respectively. Though the coefficients on the "importance" measure differs across these models (and in comparison to the PageRank model in Figure 3b) as would be expected given their different attributes, the overall interpretation is consistent across all three: At the margin, more "important" or "influential" individuals are more likely to re-locate outside of Abyan during the AQAP occupation in 2011-2012.

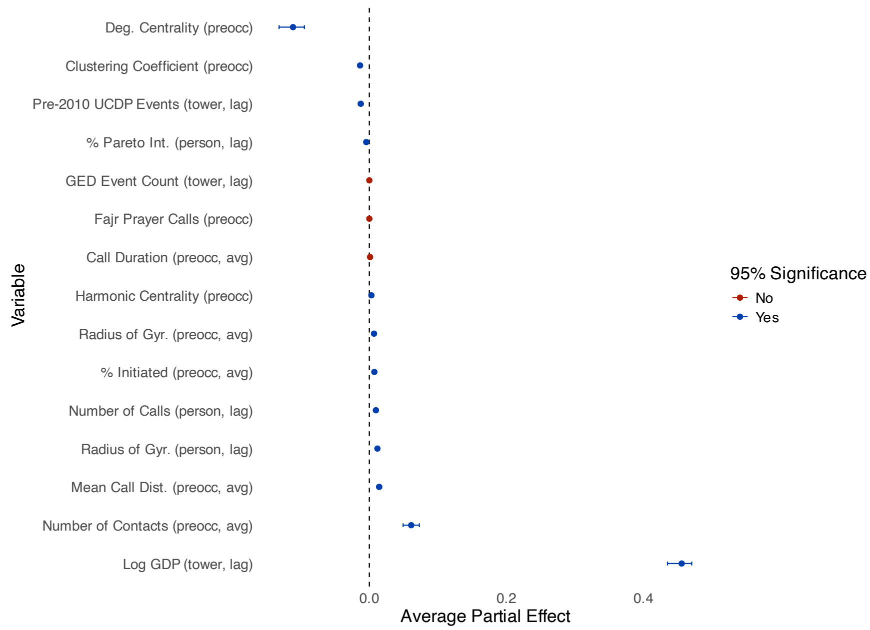

Figure A.9: Average partial effect estimates from a `stan` model re-estimating the model in Figure 3b, but using Harmonic Centrality (Marchiori and Latora, 2000; Rochat, 2009) as a measure of "influence" in place of PageRank. The substantive interpretation of the association between influence and displacement detailed in Section 6 hold with this alternative measure of social influence. Tabular results corresponding to this figure are in Table A.15.

Table A.15: Average Partial Effect estimates corresponding to Figure A.9. Signs are reversed in the plot in order to show change in the probability of leaving, not remaining

Parameter	Median Partial Effect	2.5%	97.5%
Number of Contacts (preocc, avg)	-0.0611	-0.0730	-0.0493
Call Duration (preocc, avg)	-0.0011	-0.0034	0.0011
% Initiated (preocc, avg)	-0.0073	-0.0095	-0.0053
Mean Call Dist. (preocc, avg)	-0.0143	-0.0178	-0.0115
Fajr Prayer Calls (preocc)	0.0000	0.0000	0.0000
Radius of Gyr. (preocc, avg)	-0.0069	-0.0092	-0.0052
Harmonic Centrality (preocc)	-0.0030	-0.0061	-0.0001
Clustering Coefficient (preocc)	0.0135	0.0112	0.0161
Deg. Centrality (preocc)	0.1113	0.0947	0.1316
Log GDP (tower, lag)	-0.4556	-0.4702	-0.4349
Pre-2010 UCDP Events (tower, lag)	0.0125	0.0094	0.0159
GED Event Count (tower, lag)	0.0000	0.0000	0.0000
Number of Calls (person, lag)	-0.0096	-0.0114	-0.0081
Radius of Gyr. (person, lag)	-0.0118	-0.0136	-0.0102
% Pareto Int. (person, lag)	0.0045	0.0032	0.0060

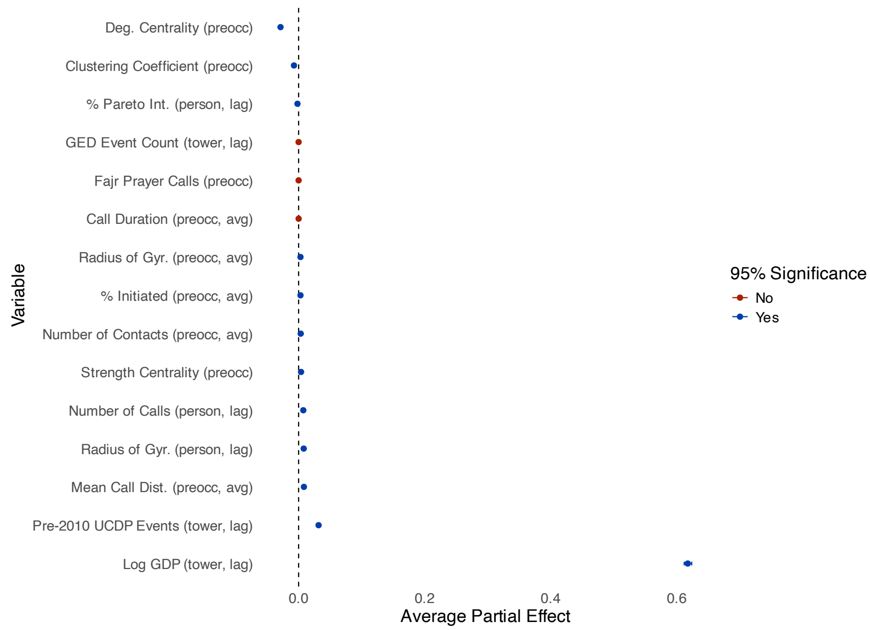

Figure A.10: Average partial effect estimates from a `stan` model re-estimating the model in Figure 3b, but using Strength Centrality (Barrat et al., 2004) as a measure of "influence" (per Opsahl et al., 2010) in place of PageRank. The substantive interpretation of the association between influence and displacement detailed in Section 6 hold with this alternative measure of social influence. Tabular results corresponding to this figure are in Table A.16

Table A.16: Average Partial Effect estimates corresponding to Figure A.10. Signs are reversed in the plot in order to show change in the probability of leaving, not remaining

Parameter	Median Partial Effect	2.5%	97.5%
Number of Contacts (preocc, avg)	-0.0033	-0.0048	-0.0019
Call Duration (preocc, avg)	-0.0000	-0.0005	0.0004
% Initiated (preocc, avg)	-0.0030	-0.0034	-0.0026
Mean Call Dist. (preocc, avg)	-0.0084	-0.0091	-0.0078
Fajr Prayer Calls (preocc)	0.0000	0.0000	0.0000
Radius of Gyr. (preocc, avg)	-0.0030	-0.0034	-0.0026
Strength Centrality (preocc)	-0.0039	-0.0045	-0.0033
Clustering Coefficient (preocc)	0.0074	0.0069	0.0079
Deg. Centrality (preocc)	0.0287	0.0268	0.0309
Log GDP (tower, lag)	-0.6173	-0.6236	-0.6118
Pre-2010 UCDP Events (tower, lag)	-0.0317	-0.0339	-0.0297
GED Event Count (tower, lag)	0.0000	0.0000	0.0000
Number of Calls (person, lag)	-0.0074	-0.0081	-0.0068
Radius of Gyr. (person, lag)	-0.0081	-0.0086	-0.0077
% Pareto Int. (person, lag)	0.0018	0.0011	0.0026

I.6 Alternative Home Detection Rules

Results presented in the body of the paper use the "modal tower" home detection rule: identifying a user's home location as the location of the cell tower to which they connect the greatest number of times in the window of interest (always a week in this paper). Recent advances in CDR methods compared in Pappalardo et al. (2021) have identified two alternate home detection rules that perform especially well for CDR data. One rule, "distinct days" identifies a user's home location as the tower to which they connect on the greatest number of days in a given window. The other "weekend modal tower" identifies a user's home location as the tower to which they connect the greatest number of times on non-work days (Thursday and Friday in Yemen as of 2011). In Figures A.11 and A.12 and Tables A.17 and A.18, we report models that use these alternative home detection rules but otherwise conform to the functional form of the main specifications reported in Figure 3b. Results support the same interpretations and conclusions, which makes sense given the density of cell towers in Yemen in 2011, as well as the coarseness of the main statistic we calculate from home location (whether a user is inside or outside Abyan governorate).

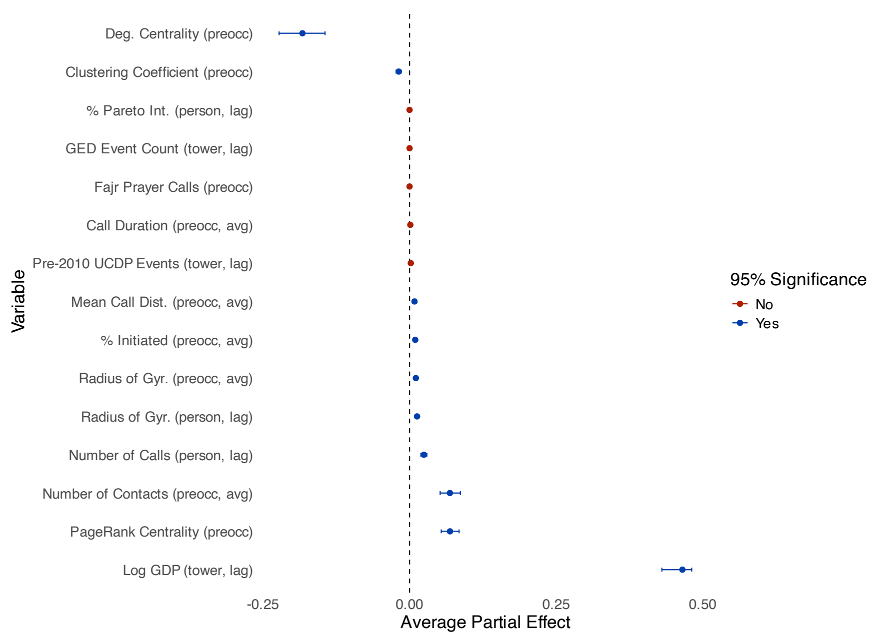

Figure A.11: Average partial effect estimates from a `stan` model re-estimating the model in Figure 3b, but calculating home location by the modal *weekend* tower, rather than the modal tower over a seven day period. This method is shown in Pappalardo et al. (2021) to perform better at identifying a user's home location. The substantive interpretation of the association between influence and displacement detailed in Section 6 hold with this alternative measure of social influence. Tabular results corresponding to this figure are in Table A.17

Table A.17: Average Partial Effect estimates corresponding to Figure A.11 Signs are reversed in the plot in order to show change in the probability of leaving, not remaining

Parameter	Median Partial Effect	2.5%	97.5%
Number of Contacts (preocc, avg)	-0.0689	-0.0867	-0.0522
Call Duration (preocc, avg)	-0.0013	-0.0041	0.0014
% Initiated (preocc, avg)	-0.0097	-0.0130	-0.0067
Mean Call Dist. (preocc, avg)	-0.0084	-0.0119	-0.0054
Fajr Prayer Calls (preocc)	0.0000	0.0000	0.0000
Radius of Gyr. (preocc, avg)	-0.0108	-0.0139	-0.0081
PageRank Centrality (preocc)	-0.0690	-0.0845	-0.0540
Clustering Coefficient (preocc)	0.0185	0.0142	0.0229
Deg. Centrality (preocc)	0.1827	0.1441	0.2225
Log GDP (tower, lag)	-0.4654	-0.4812	-0.4303
Pre-2010 UCDP Events (tower, lag)	-0.0021	-0.0053	0.0009
GED Event Count (tower, lag)	0.0000	0.0000	0.0000
Number of Calls (person, lag)	-0.0245	-0.0295	-0.0193
Radius of Gyr. (person, lag)	-0.0128	-0.0154	-0.0102
% Pareto Int. (person, lag)	0.0000	-0.0016	0.0016

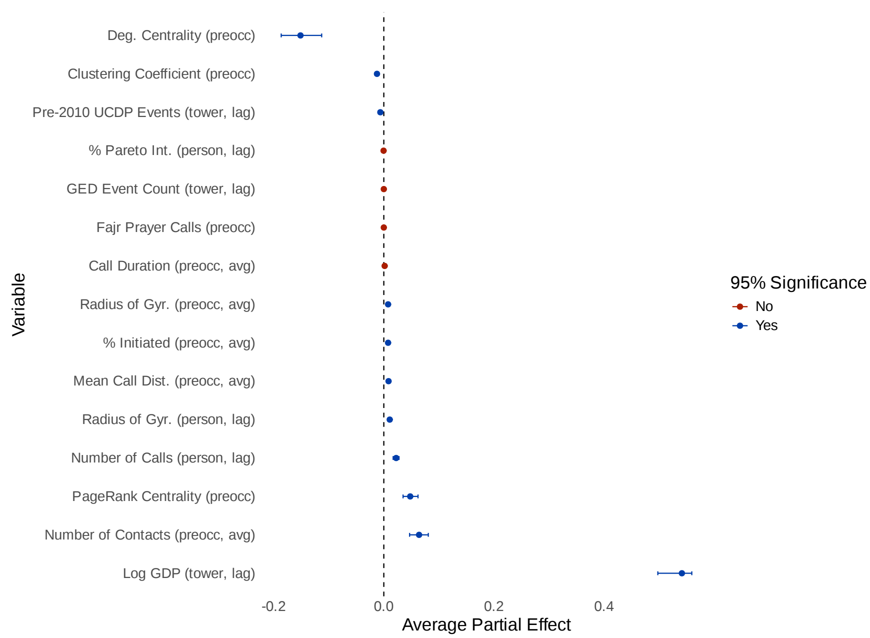

Figure A.12: Average partial effect estimates from a `stan` model re-estimating the model in Figure 3b, but calculating home location by the tower to which users connect on the greatest number of distinct days rather than the modal tower over a seven day period. This method is shown in Pappalardo et al. (2021) to perform better at identifying a user's home location. The substantive interpretation of the association between influence and displacement detailed in Section 6 hold with this alternative measure of social influence. Tabular results corresponding to this figure are in Table A.18

Table A.18: Average Partial Effect estimates corresponding to Figure A.12 Signs are reversed in the plot in order to show change in the probability of leaving, not remaining

Parameter	Median Partial Effect	2.5%	97.5%
Number of Contacts (preocc, avg)	-0.0640	-0.0808	-0.0468
Call Duration (preocc, avg)	-0.0016	-0.0040	0.0008
% Initiated (preocc, avg)	-0.0078	-0.0107	-0.0052
Mean Call Dist. (preocc, avg)	-0.0086	-0.0119	-0.0058
Fajr Prayer Calls (preocc)	0.0000	0.0000	0.0000
Radius of Gyr. (preocc, avg)	-0.0078	-0.0105	-0.0055
PageRank Centrality (preocc)	-0.0479	-0.0621	-0.0349
Clustering Coefficient (preocc)	0.0124	0.0092	0.0159
Deg. Centrality (preocc)	0.1513	0.1128	0.1863
Log GDP (tower, lag)	-0.5412	-0.5592	-0.4975
Pre-2010 UCDP Events (tower, lag)	0.0063	0.0034	0.0097
GED Event Count (tower, lag)	0.0000	0.0000	0.0000
Number of Calls (person, lag)	-0.0224	-0.0276	-0.0169
Radius of Gyr. (person, lag)	-0.0107	-0.0132	-0.0081
% Pareto Int. (person, lag)	0.0004	-0.0009	0.0018

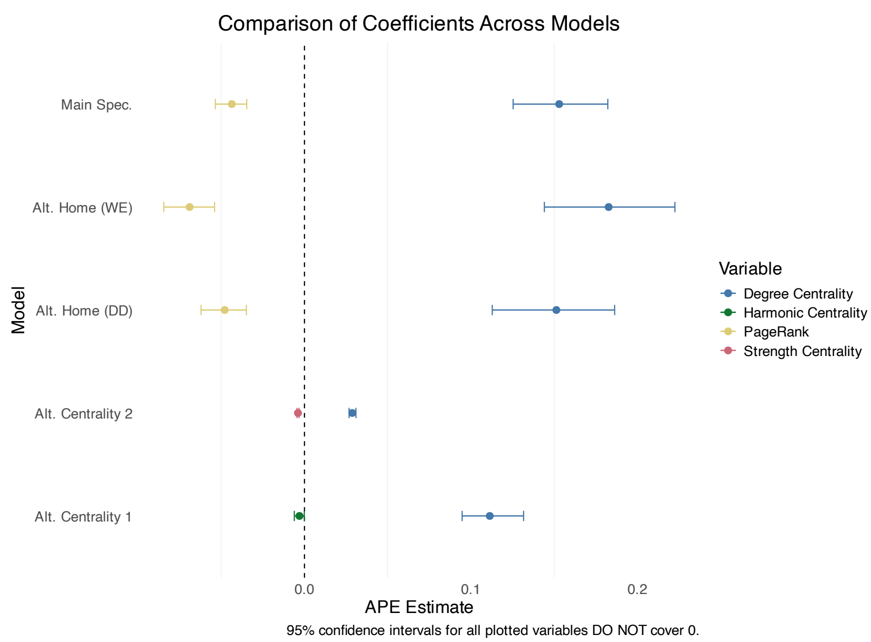

Figure A.13: Comparing the average partial effect estimates for centrality measures across different model specifications in Appendices I.5 and I.6 to the main model depicted in Figure 3b. Two models with different measures of network "importance" (Figures A.10 and A.9) and two models using different home detection rules to identify Abyan residents (Figures A.11 and A.12) all support the same substantive conclusions as the main specification. Coefficient estimates for the alternative centrality models are significantly smaller, but cannot be compared directly to the main specification because Strength centrality and Harmonic centrality are not necessarily interchangable with PageRank, nor are the population distributions (and thus the range over which the APE is calculated) necessarily similar.

I.7 Placebo Time-Period Results

Following Bagrow et al. (2011) we use pre-conflict behavior among subscribers in Abyan to assess whether our findings represent typical dynamics of migration in Abyan, or dynamics specific to the conflict period. We identify subscribers whose home location (weekly modal tower) was in Abyan during the first two weeks of January 2011 and analyze their behavior during weeks 3–12 of the pre-conflict period, ending on March 24 2011, the day before the AQAP occupation of Abyan began. We estimate the placebo model using the same functional form as the main time series results in the main text. The dependent variable is a binary indicator for whether subscribers' home location is in/out of Abyan during weeks 3–12 of the pre-conflict period. Centrality measures are computed from the same pre-occupation network graph used in the main analyses.

Table A.19 presents the full posterior estimates from `stan` estimation. The key result is a sign reversal for PageRank in the pre-conflict placebo compared to during conflict (See A.7). In non-conflict times, PageRank predicts *staying* in Abyan (posterior mean $= +1.32$, 95% CI $[1.23, 1.43]$), the opposite of the association reported in the main results from the conflict period. (posterior mean $= -0.31$, 95% CI $[-0.36, -0.26]$). Degree centrality does not reverse sign ($+0.68$ pre-conflict vs. $+1.14$ in the conflict period), suggesting it captures the migration-dampening effects of stable local social embeddedness generally, rather than a conflict-specific resource mobilization mechanism that we ascribe to PageRank.

Parameter	Mean	MCSE	SD	10%	50%	90%	n_{eff}	$\hat{R}$
# Contacts (preocc, avg)	−0.403	0.000	0.013	−0.419	−0.402	−0.386	2083	1.001
Call Duration (preocc, avg)	0.075	0.000	0.011	0.060	0.075	0.089	1677	1.001
% Initiated (preocc, avg)	−0.056	0.000	0.007	−0.065	−0.056	−0.046	2050	1.002
Radius of Gyr. (preocc, avg)	−0.054	0.000	0.006	−0.061	−0.054	−0.047	7306	1.000
% Pareto Int. (preocc)	0.026	0.000	0.009	0.014	0.026	0.037	5339	1.000
Fajr Prayer Calls (preocc)	0.012	0.000	0.008	0.002	0.012	0.023	2464	1.001
PageRank (preocc)	+1.324	0.002	0.050	1.260	1.324	1.388	736	1.004
Degree Centrality (preocc)	+0.680	0.002	0.048	0.618	0.680	0.741	632	1.007
Log GDP (tower, lag)	−0.261	0.000	0.006	−0.269	−0.261	−0.253	2662	1.001
UCDP Event Count (tower, lag)	0.340	0.000	0.009	0.329	0.340	0.352	1591	1.002
# Records (person, lag)	−0.003	0.000	0.000	−0.003	−0.003	−0.002	2244	1.002
Radius of Gyr. (person, lag)	−0.232	0.000	0.006	−0.239	−0.232	−0.225	6491	1.000
% Pareto Int. (person, lag)	0.080	0.000	0.009	0.068	0.080	0.091	5134	1.001
Latitude (tower, lag)	13.359	0.008	0.437	12.797	13.363	13.917	2840	1.002
Longitude (tower, lag)	4.762	0.003	0.144	4.578	4.763	4.946	2839	1.002

Table A.19: Posterior parameter estimates from the pre-conflict placebo model. Specification identical to the main time-series model, but fit to weeks 3–12 of the pre-conflict period. All predictors demeaned and scaled. MCSE = Monte Carlo standard error; n_{eff} = bulk effective sample size; $\hat{R}$ = potential scale reduction factor.

I.8 A Note on GDP

In Section 6, we note that the one-week lag of the GDP covariate might be mechanically inflated in the time series model if people leaving Abyan tend to a) flee to prosperous areas and b) remain in prosperous areas if they stay. This may anecdotally be true: People leaving Abyan tend to go toward Aden, which has much higher economic activity as a major port city and we know that at least some internally displaced people stayed in Aden through the end of the conflict. The lagged GDP value associated with their presence in Aden, then, is sometimes the GDP value associated with their presence in Aden—which we know to be higher than most of Abyan. Compare what we find in Figure 3b to 3a, where the GDP value is for the subscriber's starting tower *prior* to the conflict. The GDP value associated with the initial tower—a qualitatively different measure, to be clear—is not significantly associated with migration after accounting for other factors, though the direction of the insignificant association is the same.

J Stan Diagnostics

Results presented in Figures 3b, A.6, A.7 and A.8 are all estimated using the `Rstan` package, and other elements of the `mc-stan` modeling language ecosystem (The Stan Development Team, 2021). The main estimation was run the MIT Supercloud high performance computing system (Reuther et al., 2018). We used 18 cores on a single Intel Xeon Gold node with 384GB of RAM.

Additional results in the appendices were estimated using `Rstan` or the `brms` interface, relying on the same computing infrastructure at MIT. Diagnostics for those supplementary models are available in the results objects deposited as part of the replication package.

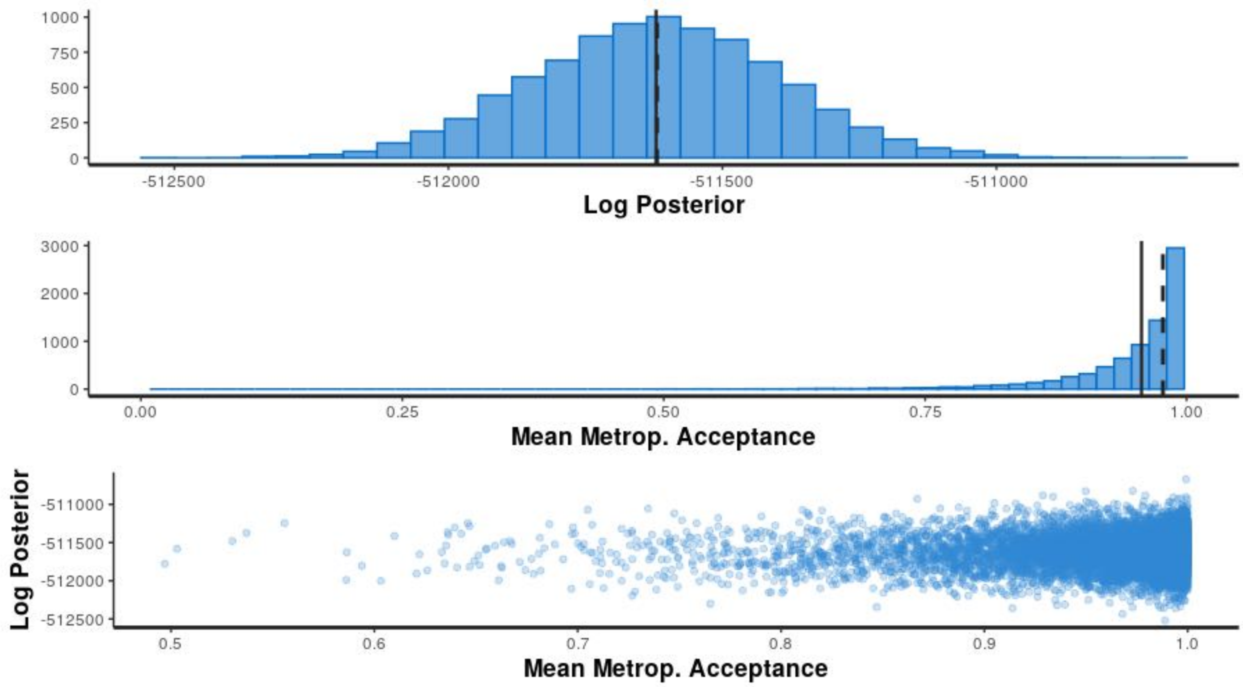

Figure A.14: Diagnostic plots from the Stan estimation we interpret in Figures A.6, 3b.

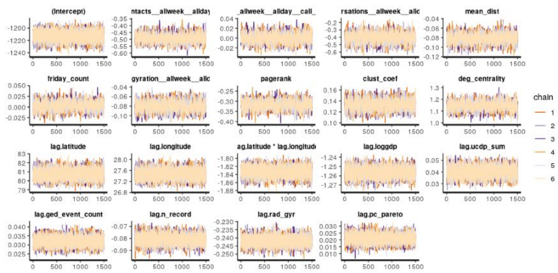

(a) Trace plots for estimation of fixed effects reported in Figures A.6, 3b.

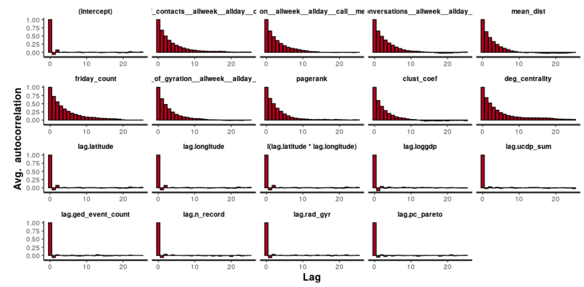

(b) Autocorrelation plots for estimation of fixed effects reported in Figures A.6, 3b.

Figure A.15: Stan diagnostics for fixed effects (coefficient estimates).

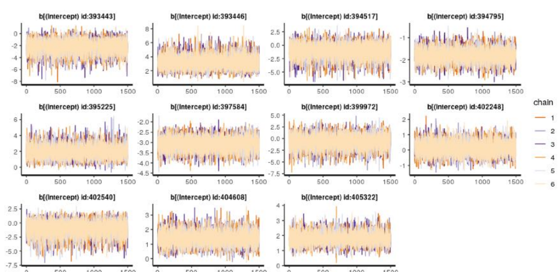

(a) Trace plots for estimation of respondent (ID) random effects from the models reported in Figures A.6, 3b.

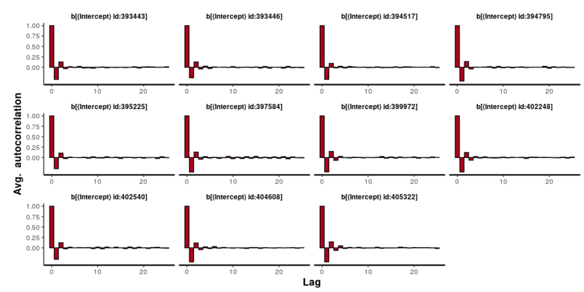

(b) Autocorrelation plots for estimation of respondent (ID) random effects from the models reported in Figures A.6, 3b.

Figure A.16: Stan diagnostics for a randomly selected chunk of respondent (ID) random effects. Respondent random effects are not interpreted, but it's good to know if they are estimated appropriately.

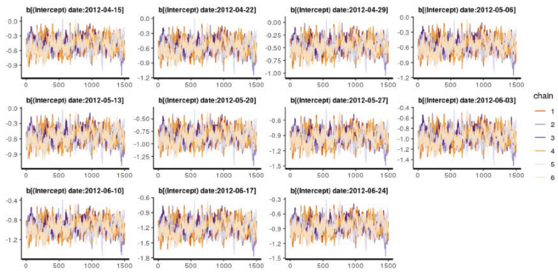

(a) Trace plots for estimation of date random effects from the models reported in Figures A.6, 3b

.

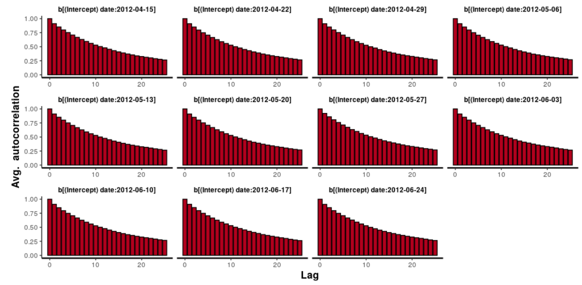

(b) Autocorrelation plots for estimation of date random effects from the models reported in Figures A.6, 3b.

Figure A.17: Stan diagnostics for a randomly selected chunk of date random effects. Date random effects are not interpreted, but it's good to know if they are estimated appropriately. The trace plots for these date parameters show good mixing, but the autocorrelation over the burn-in period is still clearly higher than for IDs and for most fixed effects. There are two reasons we are not particularly concerned about this. First, the estimates produced by this model are substantially similar to the estimates from other models (not reported here for redundancy) with different initialization. We do not interpret the date random effects, and the coefficients we do interpret do not appear to depend on one or another initialization of the date random effects. Second, per (Gelman et al., 2013, p. 284), our estimates of the date random effect parameters should not be particularly problematic. Our number of effective draws (`n_eff`) on these parameters is lower than for the ID random effects or the coefficients, but it substantially exceeds the recommended threshold of $5m$ where m is the number of chains after splitting, and the $\hat{R}$ values for the same parameters are all extremely close to the recommended value of 1.00. Estimates of these parameters may not be as precise as the fixed effects, but they do not appear problematic for interpretation of the model.

References in Supplementary Information

Abrahams, A. and D. Greenwald (2021). Geospatial research in settings of contested sovereignty. *APSA MENA Politics Newsletter 4*(1), 22–28.

Adhikari, P. (2013). Conflict-induced displacement, understanding the causes of flight. *American Journal of Political Science 57*(1), 82–89.

Alrababa'h, A., D. Masterson, M. Casalis, D. Hangartner, and J. M. Weinstein (2020). The dynamics of refugee return: Syrian refugees and their migration intentions. Working paper, 20-08, Stanford Immigration Policy Lab, Stanford, CA.

Aydoğdu, B., O. Bilgili, S. Güneş, and A. A. Salah (2025, January). Mobile phone data for anticipating displacements: practices, opportunities, and challenges. *Data & Policy 7*, e5.

Bagrow, J. P., D. Wang, and A.-L. Barabási (2011). Collective response of human populations to large-scale emergencies. *PLoS ONE 6*(3), e17680.

Balcells, L. and A. Steele (2016). Warfare, political identities, and displacement in Spain and Colombia. *Political Geography 51*, 15–29.

Barfi, B. (2010). Aqap's soft power strategy in yemen. *CTC Sentinel 3*(11), 1–5.

Barnes, T., E. Rains, J. Thomas, and J. Wu (2025). The gender digital divide and gender gaps in protest participation. Working paper, University of Texas at Austin.

Barrat, A., M. Barthélemy, R. Pastor-Satorras, and A. Vespignani (2004, March). The architecture of complex weighted networks. *Proceedings of the National Academy of Sciences of the United States of America 101*(11), 3747–3752.

Becker, S. O., V. Lindenthal, S. Mukand, and F. Waldinger (2021). Persecution and Escape: Professional Networks and High-Skilled Emigration from Nazi Germany. SoDa Laboratories Working Paper Series 2021-02, Monash University.

Berman, S. (1997). Civil society and the collapse of the weimar republic. *World Politics 49*(3), 401–29.

Bertolotti, P., A. Jadbabaie, and F. Christia (2020). Tests for network cascades via branching processes. *IEEE Transactions on Network Science and Engineering*.

Bertolotti, P., A. Milliff, F. Christia, and A. Jadbabaie (2024). Estimating the impact of drone strikes on civilians using call detail records. Technical report, Massachusetts Institute of Technology, Cambridge, MA.

Blaydes, L. and D. A. Linzer (2012). Elite competition, religiosity, and anti-americanism in the islamic world. *American Political Science Review 106*(2), 225–243.

Blumenstock, J., G. Cadamuro, and R. On (2015). Predicting poverty and wealth from mobile phone data. *Science 350*(6264), 1073–76.

Blumenstock, J., G. Chi, and X. Tan (2023). Migration and the value of social networks. *Review of Economic Studies*.

Blumenstock, J., Y. Shen, and N. Eagle (2010). A method for estimating the relationship between phone use and wealth. In *ICTD 2010*, London. ICT4 Collective.

Blumenstock, J. E. (2018, May). Estimating economic characteristics with phone data. *AEA Papers and Proceedings 108*, 72–76.

Bozcaga, T., F. Christia, E. Harwood, C. Daskalakis, and C. Papademetriou (2019). Syrian refugee integration in turkey: Evidence from call detail records. In A. A. Salah, A. Pentland, B. Lepri, and E. Letouzé (Eds.), *Guide to Mobile Data Analytics in Refugee Scenarios, The 'Data for Refugees Challenge' Study*, pp. 223–249. Springer.

Camarena, K. R., S. Claudy, J. Wang, and A. L. Wright (2020, 07). Political and environmental risks influence migration and human smuggling across the mediterranean sea. *PLOS ONE 15*(7).

Chiovelli, G., S. Michalopoulos, E. Papaioannou, and S. Sequeira (2021). Forced displacement and human capital: Evidence from separated siblings. Technical report, National Bureau of Economic Research.

Christia, F., M. Curry, C. Daskalakis, E. Demaine, J. P. Dickerson, M. Hajiaghayi, A. Hesterberg, M. Knittel, and A. Milliff (2021). Scalable equilibrium computation in multi-agent influence games on networks. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI-21)*, Palo Alto. Association for the Advancement of Artificial Intelligence.

Cruz, C., J. Labonne, and P. Querubín (2017, October). Politician family networks and electoral outcomes: Evidence from the philippines. *American Economic Review 107*(10), 3006–37.

Czaika, M. and K. Kis-Katos (2009). Civil conflict and displacement: Village-level determinants of forced migration in aceh. *Journal of Peace Research 46*(3), 399–418.

Davenport, C., W. H. Moore, and S. Poe (2003). Sometimes you just have to leave: Domestic threats and forced migration, 1964-1989. *International Interactions 29*(1), 27–55.

De Choudhury, M., W. A. Mason, J. M. Hofman, and D. J. Watts (2010, April). Inferring relevant social networks from interpersonal communication. In *Proceedings of the 19th International Conference on World Wide Web*, WWW '10, New York, NY, USA, pp. 301–310. Association for Computing Machinery.

de Montjoye, Y.-A., C. A. Hidalgo, M. Verleysen, and V. D. Blondel (2013). Unique in the crowd: The privacy bounds of human mobility. *Scientific Reports 3*(1), 1376.

de Montjoye, Y.-A., L. Rocher, and A. S. Pentland (2016). bandicoot: a python toolbox for mobile phone metadata. *Journal of Machine Learning Research 17*(175), 1–5.

Donner, J. (2007). The Rules of Beeping: Exchanging Messages Via Intentional "Missed Calls" on Mobile Phones. *Journal of Computer-Mediated Communication 13*(1), 1–22.

Dube, O., J. Blumenstock, and M. Callen (2022). Measuring religion from behavior: Climate shocks and religious adherence in afghanistan. Technical report, National Bureau of Economic Research.

Eagle, N., A. S. Pentland, and D. Lazer (2009). Inferring friendship network structure by using mobile phone data. *Proceedings of the National Academy of Sciences 106*(36), 15274–15278.

Engel, S. and A. M. Ibáñez (2007, Jan). Displacement due to violence in colombia: A household-level analysis. *Economic Development and Cultural Change 55*(2), 335–365.

Fearon, J. D. and A. Shaver (2020). Civil war violence and refugee outflows. Working paper,, Stanford University, Stanford, CA.

Felbo, B., P. Sundsøy, A. S. Pentland, S. Lehmann, and Y.-A. de Montjoye (2017). Modeling the temporal nature of human behavior for demographics prediction. Arxiv preprint, MIT Media Lab, Cambridge, MA.

Gelman, A., J. B. Carlin, H. S. Stern, D. B. Dunson, A. Vehtari, and D. B. Rubin (2013, Nov). *Bayesian Data Analysis, Third Edition*. CRC Press.

Ghosh, T., R. Powell, C. Elvidge, K. Baugh, P. Sutton, and S. Anderson (2010). Shedding light on the global distribution of economic activity. *The Open Geography Journal 3*, 147–160.

Ghoshal, G. and A.-L. Barabási (2011, July). Ranking stability and super-stable nodes in complex networks. *Nature Communications 2*(1), 394. Publisher: Nature Publishing Group.

Gonzalez, M., C. Hidalgo, and A.-L. Barabási (2008). Understanding individual human mobility patterns. *Nature 453*, 779–782.

Goodman, S., A. BenYishay, M. Lv, and D. Runfola (2019). Geoquery: Integrating hpc systems and public web-based geospatial data tools. *Computers and Geosciences 122*, 103–112.

Greenhill, K. M. (2010, Mar). *Weapons of Mass Migration: Forced Displacement, Coercion, and Foreign Policy*. Ithaca, NY: Cornell University Press.

Jampaklay, A., K. Ford, and A. Chamratrithirong (2017, Jul). How does unrest affect migration? evidence from the three southernmost provinces of thailand. *Demographic Research 37*, 25–52.

Kalyvas, S. (2006). *The Logic of Violence in Civil War*. Cambridge: Cambridge University Press.

Kalyvas, S. and L. Balcells (2010). International System and Technologies of Rebellion: How the End of the Cold War Shaped Internal Conflict. *American Political Science Review 104*(03), 415–429.

Kreindler, G. E. and Y. Miyauchi (2019, February). Measuring Commuting and Economic Activity inside Cities with Cell Phone Records. Boston University - Department of Economics - Working Papers Series WP2020-006, Boston University - Department of Economics.

Levine, L. J. and R. S. Edelstein (2010). Emotion and memory narrowing: A review and goal-relevance approach. *Cognition and Emotion*, 178–220.

Lewis, K., J. Kaufman, M. Gonzalez, A. Wimmer, and N. Christakis (2008). Tastes, ties, and time: A new social network dataset using facebook. com. *Social networks 30*(4), 330–342.

Lu, X., L. Bengtsson, and P. Holme (2012). Predictability of population displacement after the 2010 haiti earthquake. *Proceedings of the National Academy of Sciences 109*(29), 11576–11581.

Marchiori, M. and V. Latora (2000, October). Harmony in the small-world. *Physica A: Statistical Mechanics and its Applications 285*(3), 539–546.

Marston, J. (2020). Resisting displacement: Leveraging interpersonal ties to remain despite criminal violence in medellin colombia. *Comparative Political Studies 53*(13), 1995–2028.

Mironova, V., L. Mrie, and S. Whitt (2019). Risk tolerance during conflict: Evidence from aleppo, syria. *Journal of Peace Research 56*(6), 767–782.

Moore, W. H. and S. M. Shellman (2004). Fear of persecution: Forced migration, 1952-1995. *Journal of Conflict Resolution 48*(5), 723–745.

Narayanan, A. and V. Shmatikov (2009). De-anonymizing social networks. In *30th IEEE Symposium on Security and Privacy*. IEEE.

Nisbett, R. and T. D. Wilson (1977). Telling more than we can know: Verbal reports on mental processes. *Psychological Review 84*(3), 213–260.

Opsahl, T., F. Agneessens, and J. Skvoretz (2010, July). Node centrality in weighted networks: Generalizing degree and shortest paths. *Social Networks 32*(3), 245–251.

Padgett, J. F. and C. K. Ansell (1993). Robust Action and the Rise of the Medici, 1400-1434. *American Journal of Sociology 98*(6).

Page, L., S. Brin, R. Motwani, and T. Winograd (1998). The pagerank citation ranking: Bringing order to the web. Technical Report SIDL-WP-1999-0120, Stanford University InfoLab, Stanford, CA.

Pappalardo, L., L. Ferres, M. Sacasa, C. Cattuto, and L. Bravo (2021, December). Evaluation of home detection algorithms on mobile phone data using individual-level ground truth. *EPJ Data Science 10*(1), 1–19.

Pappalardo, L., E. Manley, V. Sekara, and L. Alessandretti (2023). Future directions in human mobility science. *Nature computational science 3*(7), 588–600.

Pettersson, T. and M. Oberg (2020). Organized violence, 1989-2019. *Journal of peace Research 57*(4), 597–613.

Reuther, A., J. Kepner, C. Byun, S. Samsi, W. Arcand, D. Bestor, B. Bergeron, V. Gadepally, M. Houle, M. Hubbell, et al. (2018). Interactive supercomputing on 40,000 cores for machine learning and data analysis. In *2018 IEEE High Performance extreme Computing Conference (HPEC)*, pp. 1–6. IEEE.

Revkin, M. (2019). To stay or to leave? displacement decisions during rebel governance. Working paper, Yale University, New Haven, CT.

Robinson, W. S. (1950). Ecological correlations and the behavior of individuals. *American Sociological Review 15*(3), 351–357.

Rochat, Y. (2009). Closeness Centrality Extended to Unconnected Graphs: the Harmonic Centrality Index.

Salah, A. A., A. Pentland, B. Lepri, E. Letouzé, Y.-A. de Montjoye, X. Dong, Ö. Dağdelen, and P. Vinck (2019). Introduction to the data for refugees challenge on mobility of syrian refugees in turkey. In A. A. Salah, A. Pentland, B. Lepri, and E. Letouzé (Eds.), *Guide to Mobile Data Analytics in Refugee Scenarios: The 'data for Refugees Challenge' Study*, pp. 3–27. Cham: Springer International Publishing.

Schmeidl, S. (1997). Exploring the causes of forced migration: A pooled time-series analysis, 1971-1990. *Social Science Quarterly*, 284–308.

Schon, J. (2019). Motivation and opportunity for conflict-induced migration: An analysis of syrian migration timing. *Journal of Peace Research 56*(1), 12–27.

Stopczynski, A., V. Sekara, P. Sapiezynski, A. Cuttone, M. M. Madsen, J. E. Larsen, and S. Lehmann (2014, April). Measuring Large-Scale Social Networks with High Resolution. *PLOS ONE 9*(4), e95978.

Sîrbu, A., G. Andrienko, N. Andrienko, C. Boldrini, M. Conti, F. Giannotti, R. Guidotti, S. Bertoli, J. Kim, C. I. Muntean, L. Pappalardo, A. Passarella, D. Pedreschi, L. Pollacci, F. Pratesi, and R. Sharma (2021, May). Human migration: the big data perspective. *International Journal of Data Science and Analytics 11*(4), 341–360.

The Stan Development Team (2021, February). Stan modeling language users guide and reference manual,. https://mc-stan.org.

Thompson, S. A. and C. Warzel (2019, December). Opinion | Twelve Million Phones, One Dataset, Zero Privacy. *The New York Times*.

Uppsala Conflict Data Program (2020). conflict encyclopedia. Technical report, Uppsala University.

Watts, D. J. and S. H. Strogatz (1998). Collective dynamics of 'small-world'networks. *Nature 393*(6684), 440–442.

Weiner, M. (1996). Bad neighbors, bad neighborhoods: An inquiry into the causes of refugee flows. *International Security 21*(1), 5–42.

Yabe, T., N. K. W. Jones, P. S. C. Rao, M. C. Gonzalez, and S. V. Ukkusuri (2022, June). Mobile phone location data for disasters: A review from natural hazards and epidemics. *Computers, Environment and Urban Systems 94*, 101777.

Zhao, Z., S.-L. Shaw, Y. Xu, F. Lu, J. Chen, and L. Yin (2016, September). Understanding the bias of call detail records in human mobility research. *International Journal of Geographical Information Science 30*(9), 1738–1762.

Zhukov, Y. M. (2013). Population resettlement in war: Theory and evidence from Soviet archives. *Journal of Conflict Resolution 59*(7), 1155–1185.