

# UC Berkeley

## UC Berkeley Previously Published Works

### Title

Privacy guarantees for personal mobility data in humanitarian response.

### Permalink

<https://escholarship.org/uc/item/9h7865b8>

### Journal

Scientific Reports, 14(1)

### Authors

Kohli, Nitin

Aiken, Emily

Blumenstock, Joshua

### Publication Date

2024-11-19

### DOI

10.1038/s41598-024-79561-2

Peer reviewed



## OPEN Privacy guarantees for personal mobility data in humanitarian response

Nitin Kohli<sup>1,3</sup>, Emily Aiken<sup>2,3</sup> & Joshua E. Blumenstock<sup>1,2</sup>✉

Personal mobility data from mobile phones and other sensors are increasingly used to inform policymaking during pandemics, natural disasters, and other humanitarian crises. However, even aggregated mobility traces can reveal private information about individual movements to potentially malicious actors. This paper develops and tests an approach for releasing private mobility data, which provides formal guarantees over the privacy of the underlying subjects. Specifically, we (1) introduce an algorithm for constructing differentially private mobility matrices and derive privacy and accuracy bounds on this algorithm; (2) use real-world data from mobile phone operators in Afghanistan and Rwanda to show how this algorithm can enable the use of private mobility data in two high-stakes policy decisions: pandemic response and the distribution of humanitarian aid; and (3) discuss practical decisions that need to be made when implementing this approach, such as how to optimally balance privacy and accuracy. Taken together, these results can help enable the responsible use of private mobility data in humanitarian response.

Personal mobility data has the potential to provide critical information to guide the response to humanitarian crises, and to advance several of the Sustainable Development Goals (SDGs). For instance, recent work has shown that mobility data can be used to model and prevent the spread of epidemics<sup>1–3</sup>, monitor and assist displaced populations after natural disasters<sup>4–6</sup>, target emergency cash transfers<sup>7,8</sup>, and flag riots and violence against civilians<sup>9,10</sup>. The COVID-19 pandemic in particular demonstrated that aggregated personal mobility data can provide critical insight into population movement<sup>11–15</sup>, and resulted in a push for collaborations among governments, technology companies, and researchers to leverage these data to inform evidence-based pandemic response measures<sup>16</sup>.

At the same time, personal mobility data is, by definition, personal, so its analysis – even for humanitarian purposes – may create privacy risk. Individual mobility data can expose personally identifying and sensitive information, including individuals' home and work locations, travel patterns, and interpersonal interactions. Moreover, since certain locations are endowed with social meaning, mobility data can be used to draw inferences about an individual's life, including their political preferences<sup>17</sup> and sexual orientation<sup>18,19</sup>. Even coarsening the spatial and temporal characteristics of individual locations traces provides little privacy protection, due to the idiosyncratic nature of human movements throughout daily life<sup>20</sup>. There has thus been widespread criticism and controversy about the use of such data by governments<sup>21</sup>, private companies<sup>17</sup>, civil society<sup>22</sup>, and researchers<sup>23</sup>.

To mitigate such personal inferences about individuals from digitally-derived mobility data – and the potential harms that could follow from these revelations – active measures must be taken to ensure that individual privacy is protected when such data are analyzed or released. The traditional approach to protecting privacy is to release only aggregated statistics about group movement rather than individual mobility trajectories<sup>24</sup>. However, such aggregation is typically insufficient to de-identify individuals in mobility traces<sup>20,25</sup>, so other papers have turned to other techniques including distorting spatiotemporal information<sup>26,27</sup>, obfuscating individual trajectories<sup>28</sup>, and data slicing and mixing<sup>29</sup>.

This paper develops and tests an approach to producing *provably private* mobility statistics from personal data based on differential privacy<sup>30</sup>. A number of previous studies have shown how differentially private approaches can add carefully calibrated noise to aggregated mobility statistics to ensure that the amount of individual information leaked by such statistics is bounded<sup>31–40</sup>. In comparison to past work, the focus of our study is on understanding the implications and tradeoffs that arise when using private mobility statistics in humanitarian settings. We are specifically interested in how differential privacy impacts *intervention accuracy*, i.e., the accuracy

<sup>1</sup>Center for Effective Global Action, UC Berkeley, Berkeley 94704, USA. <sup>2</sup>School of Information, UC Berkeley, Berkeley 94704, USA. <sup>3</sup>These authors contributed equally: Nitin Kohli and Emily Aiken. ✉email: jblumenstock@berkeley.edu

(and effectiveness) of policy decisions that are based on private mobility data rather than standard (non-private) data.

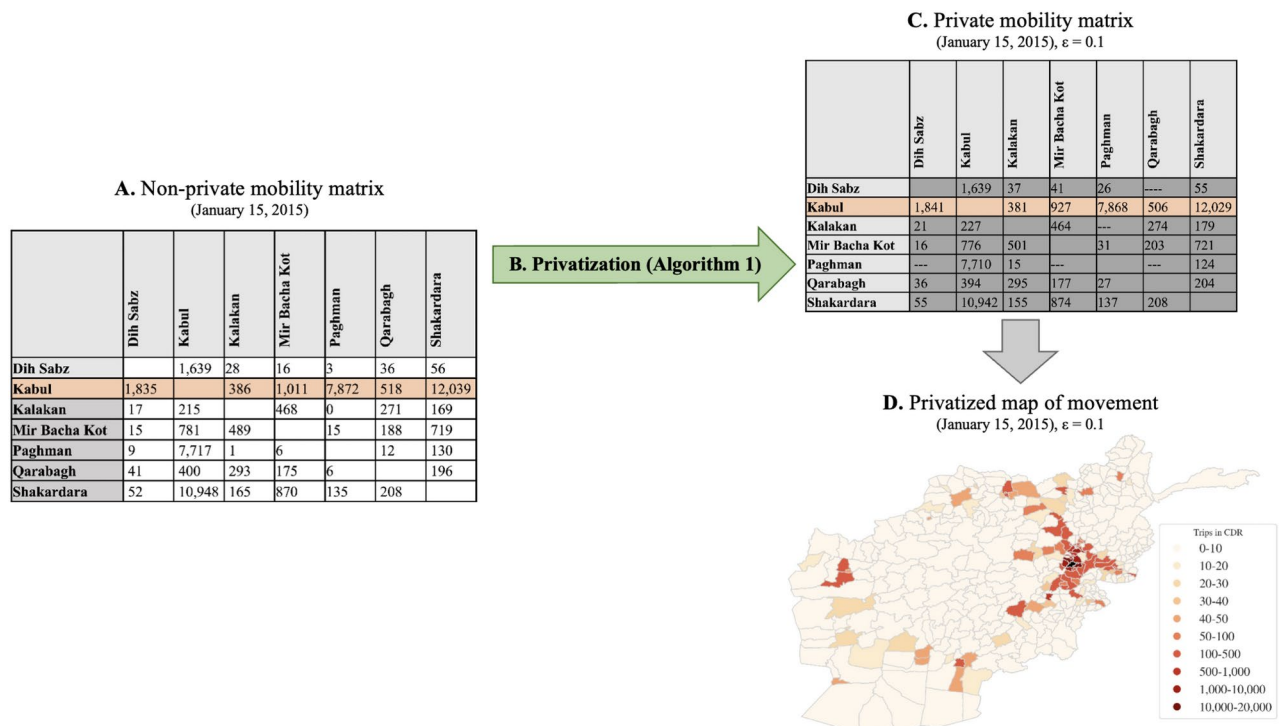
This paper makes three main contributions. First, we develop an algorithm for computing differentially private mobility information, and derive the guarantees that this algorithm provides for both data privacy and intervention accuracy. Second, we provide experimental results on the tradeoff between privacy and the accuracy of downstream policy decision in two real-world humanitarian contexts: response to (1) pandemics and (2) natural disasters and violent events. Third, we discuss the nuanced implementation choices that policymakers and algorithm developers are likely to grapple with in deploying such systems, and their effects on the privacy-accuracy tradeoff and ability to effectively deliver humanitarian aid.

## Results

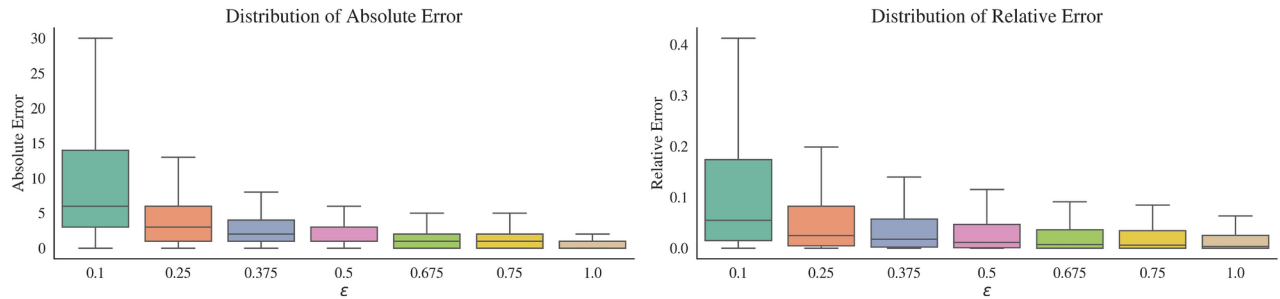
### Building and testing a differentially private mobility matrix

Our first set of results formalize the notion of a *private mobility matrix*, propose a novel algorithm for generating such matrices, and derive guarantees on the privacy and accuracy provided by this algorithm. In general, a *mobility matrix* is a data structure that quantifies the extent to which a population moves between regions in a fixed period of time—see Fig. 1A for an example of an O-D matrix derived from mobile phone records. The most common form of mobility matrix, an origin-destination matrix (hereafter abbreviated as an O-D matrix), indicates for each pair of regions ( $i, j$ ) the number of trips taken by individuals from  $i$  to  $j$ . Mobility matrices are widely used by scholars and policymakers, and are playing an increasingly critical role in many humanitarian settings<sup>1,2,13,14</sup>. However, as discussed above, the raw data in mobility matrices are sensitive, and can compromise the privacy of the individuals whose data are used to construct those matrices<sup>34,41</sup>.

Algorithm 1, described in detail in the “Methods” section, provides a computationally efficient method to construct *differentially private*<sup>30</sup> origin-destination matrices from individual mobility data. At its core, this algorithm adds very carefully calibrated noise to the underlying data, in a specific way that both protects the privacy of individuals and ensures the fidelity of the resulting O-D matrix—see Fig. 1 for a schematic diagram of this framework. In the Supplementary Information S1, we prove that our algorithm satisfies  $\epsilon$ -differential privacy (Theorem S1, Supplementary Information). This proof mathematically guarantees that no one with access to the private O-D matrix could infer “too much” about any individual trip (or any individual) whose data was contained in the matrix. The algorithm quantifies “too much” using a privacy loss parameter  $\epsilon \geq 0$  (Definition S1, Supplementary Information). This privacy loss parameter  $\epsilon$  is thus critical to understanding the



**Figure 2.** Panel A: A portion of an origin-destination matrix, showing movement between 10 districts of Afghanistan on January 1, 2015, as calculated from data provided by a mobile phone operator. Panel B: The origin-destination matrix is passed into our private O-D matrix algorithm (as described in Algorithm 1) for trip-level protection ( $T=1$ ) with privacy parameter  $\epsilon = 0.1$  and suppression threshold  $\tau = 15$ . Panel C: The same portion of the origin-destination matrix, this time after being privatized by our private O-D matrix algorithm (with  $\epsilon = 0.1$ ). Panel D: Movement from Kabul to other districts of Afghanistan on January 1, 2015, based on the highlighted row of the private origin-destination matrix in Panel B. Kabul, the origin district, is shown in black.



**Figure 3.** Left: Absolute errors in matrix entries for differentially private O-D matrices, relative to non-private matrices (derived over all days in our Afghanistan 2020 dataset and all origin-destination pairs at the admin-2 or province level). Absolute error is calculated as the difference between an O-D matrix count in the private matrix and the corresponding O-D matrix count in the non-private matrix. The distributions shown are taken over all 305 days in our 2020 CDR dataset (see Table 1) and all origin-destination province pairs. Right: Relative errors in matrix entries for differentially private O-D matrices, relative to non-private matrices. Relative error is calculated as twice the absolute error in a private O-D matrix count, divided by sum of the private and non-private O-D matrix counts for the same cell. Again, the distributions shown are taken over all 305 days in our 2020 CDR dataset (see Table 1) and all origin-destination province pairs.

	Rwanda 2008	Afghanistan 2015	Afghanistan 2020
Dates	Feb. 3–Feb. 9, 2008	Sep. 28–Oct. 4, 2015	Jan. 1–Oct. 31, 2020
Number of days	7	7	305
Number of transactions	13 million	64 million	3.2 billion
Number of subscribers	541 thousand	2.79 million	7.12 million
Admin-2 level	30 districts	34 provinces	34 provinces
Admin-3 level	416 sectors	421 districts	421 districts
5th percentile of trips per subscriber (admin-2 level)	0.00	0.00	0.00
5th percentile of trips per subscriber (admin-3 level)	0.00	0.00	0.00
Mean trips per subscriber (admin-2 level)	1.53	0.26	14
Mean trips per subscriber (admin-3 level)	2.88	1.73	52.00
95th percentile of trips per subscriber (admin-2 level)	10	2	52
95th percentile of trips per subscriber (admin-3 level)	20	9	251

**Table 1.** Summary of the call detail record datasets used in our empirical simulations.

function of the algorithm, as it governs the maximum information about an individual trip (or an individual person) is leaked by a private mobility matrix. Smaller values of  $\epsilon$  provide more robust privacy guarantees while larger values increase the risk that individual data could be compromised.

However, a central concern—and a focal point of our work—is that increases in privacy come at a cost. In particular, as  $\epsilon$  decreases, the resulting private mobility matrix looks less and less like the true (non-private) mobility matrix. In the Supplementary Information S1, we derive formal guarantees of the worst-case accuracy loss of the private O-D matrix algorithm as a function of  $\epsilon$ . In particular, Theorems 2, 3, and 4 show that the accuracy loss of our algorithm exponentially decreases as  $\epsilon$  increases; Theorem 5 further shows that, when using multiple private mobility matrices to estimate changes in population flows, the detection error decreases according to  $\Theta(\epsilon \exp(-\epsilon))$  as  $\epsilon$  increases.

To provide intuition for the tradeoff between privacy and accuracy, Fig. 2 illustrates how the privacy loss parameter  $\epsilon$  relates to differences between the original O-D matrix and the private O-D matrix. To construct the figure, we use a large dataset of mobile phone records provided by one of the largest mobile network providers in Afghanistan, covering roughly 7 million individuals and containing over 3 billion phone transactions over 305 days (for details about the data used in this paper, see Table 1). We use our algorithm to construct a private O-D matrix from these data, and compare this private matrix to a non-private O-D matrix constructed from the raw data. In the figure, we observe that the errors introduced by privatization are quite small. For  $\epsilon = 0.1$  the median absolute error (i.e., the median difference between any given private matrix count and the corresponding non-private matrix count) is 6 trips at the admin-2, and the median relative error (the absolute error divided by the non-private matrix count) is 2.54% at the admin-2 level. For  $\epsilon = 1$ —a relatively weaker level of privacy protection—the median absolute error is 0 trips and the median relative error is under 0.01% at the admin-2. As a general matter, the relative accuracy of our algorithm tends to increase as the difference between a non-private count and the suppression threshold increases; that is, lower mobility areas tend to have lower relative accuracy than higher mobility areas after applying our algorithm. This observation is consistent with prior research that

the accuracy of differentially private analyses (in general) tend to increase as the size of the underlying data increases<sup>42</sup>.

### How does privatizing mobility data affect humanitarian interventions based on those data?

The preceding results suggest that our algorithm introduces only modest errors when creating a private version of the O-D matrix. However, it is not clear if and how those modest decreases will impact downstream policy decisions based on private O-D matrices. Our next set of results therefore illustrate how policies in two key humanitarian settings are affected by the privatization of mobility data.

Specifically, we evaluate private O-D matrices in the context of a hypothetical pandemic scenario in Afghanistan, and in natural disaster scenarios in Afghanistan and Rwanda. These case studies leverage location data recorded in call detail records (CDR) collected by mobile phone operators (see Table 1)—the same sort of data that has raised privacy concerns in a wide range of contexts discussed in the Introduction<sup>20–23</sup>. The raw CDR indicate, for each phone call made on the operator's network, the date, time, and duration of the call, identifiers for each subscriber involved in each call, and the location of the cell tower through which each call was placed. This last piece of information is what makes it possible to infer the approximate location of each originating subscriber at the time of the phone call. From these raw location trace data, we calculate daily O-D matrices recording daily movements between each pair of administrative subdivisions in each country (see Fig. 1C for an example), and then show how those matrices – and the private version of them – influence downstream humanitarian interventions.

#### *Non-pharmaceutical interventions in pandemic response*

Our first case study compares the effectiveness of private and non-private mobility data in guiding public health interventions intended to stop the spread of contagious disease. This is an increasingly common application of mobility data, and an approach that was used to inform non-pharmaceutical interventions to the COVID-19 pandemic<sup>11–14</sup>. Our analysis simulates a pandemic scenario, where the spread of the disease is modeled using a classic SIR (Susceptible, Infected, Recovered) model<sup>43</sup>, adapted to account for the fact that many modern diseases are influenced by human mobility<sup>44</sup>. We consider policy decisions that may be enacted when disease prevalence estimates from the mobility-informed SIR model exceed a certain threshold (such as travel restrictions, medication distribution, and so on), and test the extent to which the timing of such decisions differs when private and non-private O-D matrices are used for mobility modeling.

To conduct these simulations, we use the same large dataset of mobile phone records from Afghanistan in 2020 used to calculate errors in matrix counts from privatization (Fig. 2). Daily O-D matrices at the province (admin-2) and district (admin-3) level derived from the CDR for 305 days are used as inputs to the mobility-informed SIR model, assuming that mobility data is available on a daily basis (but epidemiological data is available only at the start of the pandemic). We simulate three possible pandemics—pandemics initiating in Hirat, in Kabul, and in a randomly selected origin region—with 1% of individuals in the original region infected on the first day. In these simulations, we measure intervention accuracy by studying the daily (binary) decisions policymakers would make about whether or not to institute local anti-contagion policies, comparing the decisions that would be made with private data to those that would be made with non-private data. We measure the accuracy, precision, and recall of these binary decisions, with the assumption that policies are put into place at a threshold of 20% estimated disease prevalence.

As shown in Table 2, we find that intervention accuracy remains fairly high when using differentially private O-D matrices at the admin-2 (province) level (e.g. accuracy of policy decisions = 96–99% for  $\epsilon = 0.5$ ). However, when simulations are conducted with a finer spatial granularity—modeling disease spread at the district, or admin-3, level—intervention accuracy drops, and there is higher accuracy variance across pandemic simulations (e.g. accuracy of policy decisions = 72–91% for  $\epsilon = 0.5$ ). This difference highlights the importance of implementation choices like spatial granularity in deployments of the private O-D matrix algorithm: in general, the accuracy loss incurred by privatization will be larger at high spatial resolutions, when individual entries in the O-D matrix counts are smaller.

The value of the privacy loss parameter  $\epsilon$  also has important implications for intervention accuracy. Higher values of  $\epsilon$ —which correspond to less noise being added to mobility matrices in the private O-D matrix algorithm—lead to higher intervention accuracy (e.g. accuracy of policy decisions = 97–99% at the admin-2 level and 77–91% at the admin-3 level for  $\epsilon = 1.0$ ). Lower values of  $\epsilon$ —meaning more noise added to mobility matrices—lead to lower intervention accuracy (e.g. accuracy of policy decisions = 90–98% at the admin-2 level and 64–83% at the admin-3 level for  $\epsilon = 0.1$ ). The choice of  $\epsilon$  ultimately comes down to a policy decision on the relative importance of privacy and accuracy; we discuss methods for tuning the privacy-accuracy tradeoff at the end of the results section.

We also note that the precision and recall of policy decisions based on private O-D matrices can depend on the source of the pandemic, even for the same value of  $\epsilon$ . For instance, compared to Kabul, which has the highest precision and recall (Table 2), performance drops when the pandemic originates in Hirat (or in a random location). This occurs because the volume of trips is much lower in less populous regions: whereas 18,602 trips originate from Kabul on an average day, only 3,733 originate in Hirat. Thus, when we apply our algorithm, the noise and suppression introduced into small counts has a larger effect on the mobility-adapted SIR dynamics than for larger counts. This sensitivity highlights the importance of empirically evaluating the privacy-accuracy tradeoff in context, before deciding if and how to implement in policy.

#### *Geographic targeting of humanitarian aid after natural disasters and violent events*

Our second case study illustrates how the privatization of mobility data can impact the allocation of humanitarian aid (such as medical resources, temporary shelter, and food) following natural disasters and other shocks. This

	Admin-2 level (provinces)			Admin-3 level (districts)		
	Accuracy	Precision	Recall	Accuracy	Precision	Recall
Panel A: Pandemic initiating in Kabul						
Non-private	100%	100%	100%	100%	100%	100%
$\epsilon=0.1$	98% (2%)	94% (10%)	95% (10%)	83% (11%)	61% (27%)	64% (27%)
$\epsilon=0.5$	99% (0%)	98% (3%)	98% (3%)	91% (8%)	79% (21%)	80% (20%)
$\epsilon=1.0$	99% (0%)	98% (2%)	98% (2%)	91% (9%)	78% (22%)	80% (22%)
Panel B: Pandemic initiating in Hirat						
Non-private	100%	100%	100%	100%	100%	100%
$\epsilon=0.1$	90% (8%)	64% (33%)	63% (33%)	64% (12%)	16% (27%)	18% (29%)
$\epsilon=0.5$	96% (5%)	85% (21%)	85% (21%)	72% (12%)	33% (29%)	35% (30%)
$\epsilon=1.0$	97% (5%)	89% (20%)	89% (20%)	77% (11%)	44% (29%)	45% (29%)
Panel C: Pandemic initiating randomly						
Non-private	100%	100%	100%	100%	100%	100%
$\epsilon=0.1$	93% (8%)	74% (31%)	74% (32%)	71% (15%)	33% (36%)	34% (37%)
$\epsilon=0.5$	97% (5%)	89% (19%)	89% (19%)	79% (13%)	52% (31%)	53% (32%)
$\epsilon=1.0$	98% (3%)	93% (12%)	93% (12%)	82% (12%)	58% (30%)	59% (30%)

**Table 2.** Average errors in pandemic response (measured in accuracy, precision, recall) by degree of privacy introduced in the mobility matrix, at the admin-2 (left) and admin-3 level (right). Accuracy, precision, and recall are calculated for each province/district across all 305 timesteps in the simulation, and averages are taken across all provinces/districts. Standard deviations across provinces/districts are reported in parentheses.

application is motivated by the fact that mobility data derived from mobile phones and social media have been used to inform post-disaster humanitarian interventions, including following the 2010 earthquake in Haiti<sup>4</sup>, the 2019 wildfires in California<sup>45</sup>, and most recently during the 2023 earthquake in Turkey<sup>46</sup>. In our hypothetical aid-targeting scenario we simulate a short-term shock—such as an earthquake, landslide, flood, or violent event—and measure the degree of out-migration (as the total out-trips observed in an O-D matrix) in the following week. We assume a government or other policy actor seeks to provide humanitarian aid to the migrants in the areas with the most out-migration.

For our aid targeting simulations, we use two CDR datasets related to real-world natural disasters and violent events in simulating the implications of our private O-D matrix algorithm for targeting humanitarian aid. In Afghanistan, we study the battle of Kunduz in 2015, which resulted in the displacement of more than 100,000 people according to official statistics<sup>47</sup>. We focus on displacement in the week following the largest single day of the battle, using a CDR dataset of around 64 million transactions associated with 2.8 million mobile subscribers (Table 1). In Rwanda, we study mobility in the week following the Lake Kivu Earthquake of 2008, using a CDR dataset that includes around 13 million transactions associated with half a million mobile subscribers (Table 1). In these settings we measure intervention accuracy using two metrics: first, we compare estimated total out-migration from the affected area using private data to non-private data, and second we calculate the accuracy of the top 3 regions of out-migration (i.e., the regions receiving the most migrants after a disaster) identified in private data in comparison to those identified in non-private data.

In Table 3 we show that in both the Afghanistan and Rwanda contexts, private O-D matrices with trip-level protection have relatively low error in determining total out-migration following natural disasters and violent events in comparison to non-private data (e.g. percent error = 2.54–8.27% for  $\epsilon = 0.5$ ). All private counts are biased downwards relative to non-private counts due to suppression of small values in private O-D matrices (Fig. S2, Supplementary Information). As in the simulations of non-pharmaceutical interventions in pandemic response, we observe that intervention accuracy using private mobility matrices is higher at the admin-2 level (percent error = 2.54–7.38% for  $\epsilon = 0.5$ ) than at the admin-3 level (percent error = 4.98–8.27% for  $\epsilon = 0.5$ ). Unintuitively, however, we actually observe *higher* intervention accuracy for small values of  $\epsilon$  (which correspond to large amounts of noise added to mobility matrices) than for large values of  $\epsilon$  (which correspond to adding small amounts of noise). This arises from the use of a suppression threshold. By Theorem S2 (Supplementary Information), a non-private count below the threshold is more likely to become non-suppressed in Step 3 of the algorithm when a smaller  $\epsilon$  is used compared to a larger one, thereby making the total out-migration flow more accurate for  $\epsilon = 0.1$  compared to  $\epsilon \in \{0.5, 1\}$ .

Private O-D matrices with trip-level protection also have relatively low error in identifying the areas of largest out-migration for targeting humanitarian aid. We focus in our simulations on selecting the top 3 geographies with the most out-migration for targeting post-disaster resources. We find that selection with private data is 90–100% accurate for  $\epsilon = 0.5$  (Table 3). Unlike in measuring total out-migration, we observe little sensitivity to the spatial granularity of simulation or the value of  $\epsilon$  in identifying the top regions of out-migration. This result may be at least partially explained by the fact that identifying the top regions of out-migration from a relatively small number of regions with large out-migration flows is a less nuanced task than identifying the exact degree of outflow. In general we also do not observe much sensitivity to the number of regions targeted for humanitarian

	Admin-2 level				Admin-3 level			
	Non-Private	$\epsilon = 0.1$	$\epsilon = 0.5$	$\epsilon = 1$	Non-Private	$\epsilon = 0.1$	$\epsilon = 0.5$	$\epsilon = 1$
Panel A: Battle of Kunduz in Afghanistan								
Total out-migration	49,994	49,001	48,725	48,722	87,007	83,327	82,671	82,609
Percent error in total out-migration	0.00%	1.98%	2.54%	2.61%	0.00%	4.23%	4.98%	5.05%
Accuracy of top- $k$ regions ( $k = 3$ )	0.00%	90.47%	90.47%	90.47%	0.00%	85.71%	90.47%	90.47%
Panel B: Lake Kivu Earthquake in Rwanda								
Total out-migration	32,627	30,595	29,930	29,801	51,102	48,085	47,331	47,320
Percent error in total out-migration	0.00%	6.23%	8.27%	8.66%	0.00%	5.90%	7.38%	8.66%
Accuracy of top- $k$ regions ( $k = 3$ )	100.00%	90.47%	95.24%	90.47%	100.00%	100.00%	100.00%	100.00%

**Table 3.** Average errors in total migration counts and top- $k$  regions of out-migration, by degree of privacy introduced in the mobility matrix, at the admin-2 level (left) and admin-3 level (right).

aid (for example, if 5 or 10 regions are identified instead of 3, as shown in Fig. S3 (Supplementary Information S1)).

Private O-D matrices with individual-level protection exhibit similar behavior as those with trip-level protection (as specified by the parameter  $T$  in our algorithm). Table S3 (Supplementary Information) shows that as increasingly restrictive values of  $T$  are used, our total out-migration measure under-counts total out-migration. However, the value of  $T$  does not have much impact on the accuracy of identifying the top- $k$  districts of out-migration districts. Compared to trip-level privacy—which provides weaker privacy protections to individuals—we find that the privacy-accuracy tradeoff between individual- and trip-level privacy to be nuanced. In some situations, the accuracy of out-migration using individual-level protection is higher than using trip-level privacy (e.g., for the Lake Kivu earthquake at the admin-2 level, matrices with individual-level protection experience 4.87–5% error (Table S3, Panel D), compared to 6.23–8.88% error at the trip-level (Table 3, Panel B)); in other situations, the relationship is reversed (e.g., for the Battle of Kunduz at the admin-2 level, matrices with individual-level protection incur 15.39–18.15% error (Table S3, Panel A), compared to 1.98–2.61% error at the trip-level (Table 3, Panel A)). We observe a similar phenomenon in identifying the top- $k$  districts: in some situations, private O-D matrices with individual-level protection are more accurate at identifying the top- $k$  districts of out-migration (e.g., for the Lake Kivu earthquake at the admin-2 level, matrices with individual-level protection are 100% accurate (Table S3, Panels D–F), whereas matrices with trip-level protection are 90.47–95.25% accurate (Table 3, Panel B)); in other circumstances, private matrices with trip-level protection are more accurate (e.g., for the Battle of Kunduz at the admin-2 level, matrices with individual-level protection are 85.71% accurate (Table S3, Panels A–C), whereas matrices with trip-level protection are 90.47% accurate (Table 3, Panel A)).

### How should policymakers navigate the tradeoff between privacy and accuracy?

Our main results show that policymakers can switch from non-private to private mobility matrices—which guarantee a certain degree of protection of individual mobility traces—with relative little accuracy loss. For example, in simulating the accuracy of nonpharmaceutical interventions in mobility-informed pandemic response, we found an accuracy of 72–99% for policies enacted using private mobility matrices. Accuracy is similar for policy decisions related to mobility-informed targeting of humanitarian resources following natural disasters, at 77–99%. However, our results also underscore that there is a key tradeoff between privacy and accuracy (shown in Fig. 2 and made explicit in Theorems S1–S5, Supplementary Information). In the case of the private O-D matrix algorithm, this tradeoff is quantified by the privacy loss parameter  $\epsilon$ .

What is the benefit of increasing privacy by reducing  $\epsilon$ ? So far, we have shown that decreasing  $\epsilon$  generally comes with a decrease in accuracy for downstream policy interventions. All else being equal, we would like to maintain the highest standard of accuracy in policy interventions informed by mobility data—but doing so comes at a loss to the privacy of individual mobility traces. In this section, we use properties of differential privacy (Lemmas S2 and S3, Supplementary Information) to derive worst-case privacy guarantees for the policy intervention simulations from the previous section. We then compare these worst-case privacy guarantees to the expected individual-level privacy loss, and discuss practical methods for policymakers to determine the appropriate tradeoff between privacy and accuracy.

#### Worst case vs. expected privacy loss

The privacy guarantees for our differentially private O-D matrix algorithm (Lemmas S2 and S3, Supplementary Information) suggest that worst-case privacy loss for individuals can be quite high. Privacy loss is a unitless quantity that measures the maximum additional information an adversary learns about any individual when their data is used in our private O-D matrices - such information includes whether an individual traveled at least once, where they traveled to and from, and any idiosyncratic travel patterns and routines they may have, among others. As the number of days used in the simulation increases, this releases additional O-D matrices that adversaries can use to infer personal information about individuals. The mathematics of  $\epsilon$ -differential privacy formalize this intuition, and indicates that the likelihood an individual's data is compromised increases linearly in the amount of information they contribute. For example, in the pandemic simulation in Afghanistan, 305 private mobility matrices are released. When trip-level privacy is used, the privacy loss for an individual over the

entire simulation is  $\epsilon(n_1 + \dots + n_{305})$ , where  $n_j$  is the number of inter-region trips an individual took on day  $j$  that appear in the CDR. Likewise, when trip-level privacy is used in our case studies of the Lake Kivu Earthquake in Rwanda and the Battle of Kunduz in Afghanistan, where 7 days of O-D matrices are released, the privacy loss over the entire simulation is  $\epsilon(n_1 + \dots + n_7)$ , where  $n_j$  is the number of inter-region trips an individual took on day  $j$  that appear in the CDR. In both cases, the worst-case privacy loss is defined as the maximum privacy loss an individual incurs in a simulation. However, our empirical results indicate that expected privacy loss is substantially lower.

In both case studies, the expected privacy loss across the population is much lower than the worst-case characterization. In our pandemic case study in Afghanistan, the average individual makes a total of  $n_1 + \dots + n_{305} = 14$  province-to-province trips and  $n_1 + \dots + n_{305} = 52$  district-to-district trips during the 305 days of the simulation; these trips are aggregated in mobility matrices (Table 1), so the average individual-level privacy loss is  $14\epsilon$  at the admin-2 level and  $52\epsilon$  at the admin-3 level for the entirety of the simulation. For a value of  $\epsilon = 0.5$ , this corresponds to a total privacy loss of 7 for provinces and 26 for districts; the privacy loss accordingly shrinks with smaller values of  $\epsilon$ . This means that, using the 305 released private O-D matrices, the information an adversary gains about an individual (such as whether they traveled, and if so, to and from where) is, on average, at most 7 times higher at the province level (and at most 26 times higher at the district level) than had their data not been used at all. (For comparison, in the non-private approach, the information an adversary gains about the average individual is, in principle, unbounded; but using differential privacy, privacy loss can be limited by the choice of  $\epsilon$  and the number of O-D matrices released.) In our case study of the Lake Kivu earthquake in Rwanda, only seven days of mobility matrices are released, and the average individual makes 1.53 trips at an admin-2 level and 2.88 trips at an admin-3 level (Table 1). For a value of  $\epsilon = 0.5$ , this corresponds to a total privacy loss of 0.77 for provinces and 1.44 for districts. Privacy loss in our case study of the Battle of Kunduz in Afghanistan is expected to be even smaller: over the seven days of data released, the average individual makes 0.26 trips at an admin-2 level and 1.73 trips at an admin-3 level (Table 1). For a value of  $\epsilon = 0.5$ , this corresponds to a total privacy loss of 0.13 for provinces and 0.87 for districts.

While there are few benchmarks against which we can compare the expected privacy loss of our algorithm, the numbers we observe are within the range of privacy losses observed in existing differential privacy deployments. For instance, Google's COVID-19 Community Mobility Reports incurs a total privacy loss of 2.64 per day<sup>33</sup>. While Google users are likely more mobile than the individuals in our dataset, this provides a rough point of comparison. If Google's COVID-19 Community Mobility Reports were released for seven days, as in our disaster response scenarios, the total privacy loss would accumulate to 18.48. If Google's COVID-19 Community Mobility Reports were released for 305 days, as in our pandemic response scenario, the total privacy loss would accumulate to 805.2.

To further stress-test our method, we implement a state-of-the-art membership inference attack from Pyrgelis et al. (2017)<sup>48</sup>, which uses machine learning tools to attempt to predict whether an individual's information is included in an aggregate dataset (see Supplementary Information S1). We find that the raw (non-private) mobility matrices we test do indeed leak information about membership of individuals in the construction of the matrix: machine learning classification models achieve an average AUC of 0.62–0.63 in a held-out test set (Table S1). The risk of membership inference implied by these AUC values is low in absolute terms and in comparison to prior work<sup>48</sup>, but nonetheless indicates that attackers could gain access to individual information through the non-private mobility matrix. When our method is used to privatize the mobility matrix prior to the membership inference attack, the attacker's accuracy is lowered in some, but not all, cases (AUC decrease of 0–3% in Table S1). These modest decreases in AUC can have meaningful implications at scale: for instance, roughly 100,000 fewer individuals could be identified as being more likely to exist in the O-D matrix than a random individual (of the 7.12 million individuals in the entire dataset). The protections provided by our method are less than those of the empirical tests of differential privacy in Pyrgelis et al. (2017), in part because the baseline leakage of individual information (absent privatization) is much lower. Our setting is also challenging because we simulate providing the adversary with considerably more data per individual than in prior work<sup>48</sup>.

#### *Methods to tune the privacy-accuracy tradeoff*

How should policymakers choose a value of  $\epsilon$  to navigate the privacy-accuracy tradeoff? Across applications of differential privacy, there is no consensus on the “right” way to set  $\epsilon$ <sup>42,49–51</sup>. In humanitarian applications, where accurate statistics can directly improve individual and social welfare, one realistic approach is to maximize the amount of privacy an algorithm can provide subject to accuracy constraints that are determined by programmatic requirements<sup>52,53</sup>.

In situations where policymakers have a maximum tolerance for accuracy loss, one might imagine using estimates of intervention accuracy loss from the datasets being released to set  $\epsilon$ . For instance, consider an aid organization seeking to target relief to the top-3 Admin-3 locations with the highest out-migration during the Battle of Kunduz. Suppose the aid organization wants to provide the strongest privacy protections possible, provided the algorithm has at least 90% intervention accuracy. The aid organization generates Table 3 and chooses the smallest value of  $\epsilon$  that yields 90% accuracy, namely  $\epsilon = 0.5$ . Then, the aid organization uses our algorithm with  $\epsilon = 0.5$  to generate private O-D matrices to determine the 3 locations to receive relief. (Note that this approach, while intuitively appealing, has a subtle caveat: technically, the parameters of an algorithm, such as the value of  $\epsilon$ , can leak privacy when chosen based on the data<sup>51,54–56</sup>. One approach to address this would be to compute private mobility matrices on pre-program data and set  $\epsilon$  based on simulations on those data, before applying the chosen  $\epsilon$  to the deployment data).

Alternatively, policymakers could set  $\epsilon$  using heuristics based on our algorithm<sup>57</sup>. For example, one simple heuristic approach is to determine  $\epsilon$  based on a policymaker's tolerance for error of an O-D matrix count based on the noise introduced by Step 2 of Algorithm 1. The standard deviation of the Laplace distribution with



$\lambda = \epsilon^{-1}T$  is given by  $T\sqrt{2}/\epsilon$ , so we can bound on the set of feasible  $\epsilon$  for a given O-D matrix count accuracy  $\alpha$ , as  $\alpha \geq \sqrt{2}T/\epsilon \iff \epsilon \geq \sqrt{2}T/\alpha$ . For example, if policymakers require that the typical error in each entry of the O-D matrix is at most 10 (trip-level privacy), then they can set  $\epsilon \geq \sqrt{2}/10 \approx 0.14$ . However, if policymakers instead can tolerate a typical error of at most 50, then any  $\epsilon \geq \sqrt{2}/50 \approx 0.028$  will suffice. While simple, this heuristic has two drawbacks: First, it ignores the impact of the post-processing steps of our algorithm, which can affect the overall accuracy of our algorithm<sup>58–60</sup> and downstream applications<sup>61–63</sup>; And second, this heuristic only utilizes the first two moments of the Laplace distribution to find a rough balance between the privacy level  $\epsilon$  and the typical error.

A final option is to use the privacy-accuracy tradeoff theorems provided in the Supplementary Information to configure the algorithm to behave optimally. Theorems S6 and S7 (Supplementary Information) provide explicit formulas for policymakers to set  $\epsilon$  based on context-specific needs. For example, suppose policymakers require 95% confidence that the error in an entry of the private O-D matrix be at most 10. Using Theorem 6, the policymakers can utilize any  $\epsilon \geq -(10.5)^{-1} \ln(0.05) \approx 0.285$ . Compared to the value of 0.14 from the simple heuristic above, Theorem 6 enables policymakers to achieve a similar goal of bounding errors in each cell by 10 while utilizing a larger  $\epsilon$ , thereby enabling the construction of more accurate private matrices. Theorem 7 can be similarly utilized to readily translate a policymaker's tolerance for error in terms of the parameter  $\epsilon$ .

In addition to  $\epsilon$ , the choice of the suppression threshold  $\tau$  can also affect the privacy-accuracy tradeoff. As is the case with  $\epsilon$ , there is no single way to set the suppression parameter. Historically, different organizations have used varying suppression thresholds based on their perception of privacy risk. For example, the US Centers for Disease and Control and Prevention use a suppression threshold of 15 for their cancer statistics<sup>64</sup>, whereas California's Department of Education uses a threshold of 10 for their public education statistics<sup>65</sup>. In situations where an organization already has a suppression standard  $s$ , it can be used to set  $\tau$  based on the noise introduced by our algorithm. (Since the threshold only applies to Step 3 of the algorithm, it is sufficient to consider the impact of the noise from Step 2 of the algorithm). We consider two natural heuristics. The first heuristic leaves the suppression unchanged ( $\tau = s$ ), as the expected noise introduced by the Laplace distribution is 0. While this heuristic allows  $\tau$  to account for the average noise introduced, it does not capture the variability of the noise introduced by the Laplace distribution. Thus, a second heuristic sets  $\tau = s \pm \sqrt{2}/\epsilon$ , where the + or - is chosen based on domain-specific concerns. For example, if the policymaker's intervention accuracy requires few cells to be suppressed, they can set  $\tau = s - \sqrt{2}/\epsilon$  and provide additional privacy protections to individuals via data minimization. Alternatively, if the policymaker is concerned that small sub-populations could now appear in the private O-D matrix due to the introduction of noise, they could set  $\tau = s + \sqrt{2}/\epsilon$  to reduce the likelihood of such an event.

## Discussion

In this paper, we presented a differentially private algorithm for releasing O-D matrices to learn about the mobility of a population without learning “too much” about any individual (or any individual trip). We analytically derived closed-form expressions for the privacy-accuracy tradeoff and showed that our algorithm can be configured to produce accurate private mobility matrices for a wide-range of parameter settings.

We then tested the performance of our differentially private O-D matrix algorithm using mobility data from two humanitarian settings. In these settings, we compare policy decisions made using the original mobility data to decisions made from a privatized version of our data. This allows us to empirically characterize the privacy-accuracy tradeoff, and calibrate privacy in a way that allows the policymaker to still make effective decisions. Our closing discussion provides more general and practical guidance for calibrating privacy (i.e., setting  $\epsilon$ ) based on the policymaker's tolerance for uncertainty.

As digital devices proliferate throughout the developing world, the data generated by those devices are increasingly being used to advance the Sustainable Development Goals. In these and other “social good” applications, it is imperative to consider and respect the privacy of the individuals behind the data. Our algorithm, analysis, and discussion highlights one way to provide strong privacy guarantees while still allowing for policymakers to make informed decisions. We hope future work can improve and expand upon these methods, to provide a more robust set of options for using private data in effective humanitarian response.

## Methods

### Differentially private O-D matrices

The main text provides an intuitive description of the differentially private O-D matrix algorithm. Here we formalize the algorithm, and describe a number of accuracy and privacy guarantees. We defer the formal presentation of these statements, as well as their proofs, to the Supplementary Information S1.

Our Private O-D Matrix algorithm (Algorithm 1) requires four inputs: the dataset  $d$  which is used to construct the private O-D matrix, a value of  $\epsilon$  to control the privacy-accuracy tradeoff in the private matrix, a value  $T$  which has two interpretations based on whether trip-level differential privacy or individual-level differential privacy is required, and a threshold parameter  $\tau \in \mathbb{Z}_{\geq 0}$  that ensures the noisy matrix entries are non-negative and enables suppression of small counts when desired. We describe these last two parameters in more detail below.

When individual-level differential privacy is required,  $T$  represents the maximum number of trips that any one individual can contribute to the mobility matrix. In order to ensure that the dataset  $d$  does not contain more than  $T$  trips for any individual, a preprocessing transformation can be applied to censor additional trips. One such algorithm that does so runs as follows. For each individual, the algorithm determines the number of trips  $\theta$  they took. If  $\theta > T$ , the algorithm selects  $T$  trips uniformly at random to keep and drops the remaining  $\theta - T$  trips from the dataset. If  $\theta \leq T$ , the algorithm leaves all  $\theta$  trips in  $d$ . Alternatively, when trip-level differential

privacy is required, setting  $T = 1$  achieves this requirement, as the presence or absence of any trip changes any count in the O-D matrix by 1.

For any choice of  $\tau \in \mathbb{Z}_{\geq 0}$ , the entries of the private matrix are guaranteed to be non-negative. The magnitude of  $\tau$  determines the level of suppression present in the private output. When  $\tau = 0$ , there is no suppression of the values in the matrix. As we increase  $\tau$ , we increase the set of small counts to be suppressed. Prior to provable privacy techniques, it was common practice in statistical disclosure control to suppress information derived from a small number of individuals<sup>66</sup>. While suppression alone offers no provable privacy guarantee, it can inform end-users that the statistic may be unreliable due to the limitations in the data collection methods, such as having low phone coverage in a particular region. Additionally, for certain types of humanitarian responses – such as those involving disease spread – regions with small mobility counts are likely to provide fewer vectors for disease spread (and hence could be less relevant for mitigation strategies) compared to those regions with higher mobility counts. Consequently, suppression may be a useful signalling apparatus to policymakers.

---

**Input:**

- Dataset  $d$
- Privacy parameter  $\epsilon \geq 0$
- Trip threshold  $T \in \mathbb{Z}_{\geq 1}$ : For individual-level differential privacy,  $T$  corresponds to the maximum number of trips an individual can contribute to the non-private O-D matrix; the dataset  $d$  should be preprocessed to ensure every individual contributes at most  $T$  trips. For trip-level differential privacy, set  $T = 1$ ; no additional preprocessing of  $d$  is required.
- Suppression threshold  $\tau \in \mathbb{Z}_{\geq 0}$

**Output:**  $\epsilon$ -differentially private O-D matrix  $\hat{M}(d)$

- 1 (*Non-private computation*) For a geographic region with  $k$  zones, compute the  $k \times k$  matrix  $M(d)$ , where  $M(d)_{a,b}$  equals the number of trips from zone  $a$  to zone  $b$  when  $a \neq b$ , and 0 when  $a = b$ .
  - 2 (*Incorporate privacy*) Add independent Laplace noise with scale  $\lambda = \epsilon^{-1}T$  to each non-diagonal entry of  $M(d)$ . Call this resulting matrix  $M'(d)$ .
  - 3 (*Post-processing*) Next,
    - (I) Round each of the resulting entries in  $M'(d)$  to the nearest integer. In the event where there is no unique nearest integer, we adopt the convention of rounding up to the nearest integer.
    - (II) Then, for all resulting values that are below a threshold of  $\tau$ , map these values to 0.
 Call this resulting matrix  $\hat{M}(d)$
  - 4 **Return**  $\hat{M}(d)$
- 

**Algorithm 1.** Private O-D Matrix

---

This algorithm is accompanied with a several provable guarantees. For starters, our algorithm satisfies  $\epsilon$ -differential privacy (Theorem S1, Supplementary Information). In addition to the privacy guarantee, our algorithm has strong accuracy guarantees, which we call the *privacy-accuracy tradeoff theorems*. With high probability: (1) the algorithm preserves cells that would have been suppressed absent differential privacy; (2) the algorithm does not suppress cells that would not have been suppressed absent differential privacy; (3) for non-suppressed cells, the algorithm's error decays exponentially as  $\epsilon$  increases; and (4) for non-suppressed cells, when examining the difference in cell counts across different time periods, the algorithm's error decreases according to  $\Theta(\epsilon \exp(-\epsilon))$  as  $\epsilon$  increases. These statements are made rigorous in Theorems S2, S3, S4, and S5 respectively in the Supplementary Information. Theorems 6 and 7 utilize Theorems 4 and 5 to provide policymakers with exact formulas to calculate  $\epsilon$  based on their context-specific needs.

## Empirical simulations

### Data

Our simulations that use the private O-D matrix algorithm for real-world policy decisions rely on location traces derived from mobile phone metadata (call detail records, or CDR). We use three call detail record datasets: one from Afghanistan in 2015, one from Afghanistan in 2020, and one from Rwanda in 2008. Table 1 summarizes these data. Our dataset from Afghanistan in 2020 is 305 days long and contains transactions from around 7 million subscribers; the remaining two datasets are only 7 days long (as the focus on specific natural disasters or violent events) and contain data for half a million (Rwanda) and around three million (Afghanistan) subscribers, respectively. In each dataset the CDR include the date, time, and duration of each call placed on the mobile phone networks, along with pseudonymized IDs for the originating subscriber and the recipient. They also record the cell tower through which each call was placed, providing a measure of the approximate location of each originating subscriber. We derive daily O-D matrices at the admin-2 and admin-3 level for each dataset by geolocating cell towers to administrative subdivisions, and then counting the number of trips observed in the data on a daily basis (where a trip occurs if a subscriber places a call or text in one administrative subdivision and a subsequent call or text in a different subdivision).

### Algorithm parameters

Across all simulations, we use a suppression threshold of 15, which is used in practice by the US Centers for Disease and Control and Prevention in their cancer statistics<sup>64</sup>, as well as the US Department of Health and Human Services National Institute on Minority Health and Health Disparities<sup>67</sup>. For information on the share of O-D matrix counts that are suppressed in each dataset, see Table S2 in Supplemental Information. We set  $T = 1$  to provide trip-level protections with  $\epsilon \in \{0.1, 0.5, 1\}$ . We also vary the value of  $T$  in the geographic targeting simulations to quantify the impact that individual-level privacy has on the utility of downstream policy decisions (described below).

### Non-pharmaceutical interventions in pandemic response

For this case study, we rely on the mobility-based SIR (Susceptible, Infected, and Recovered) model introduced by Goel and Sharma<sup>44</sup>, which adapts the classic SIR model of Kermack and McKendrick<sup>43</sup> by including features to model human mobility. In the mobility-based SIR model, at any time  $t$ , individuals in region  $i \in [k]$  are categorized into one of three groups: they are either susceptible to the disease, infected with the disease, or recovered from the disease. The number of individuals in these groups is given by  $S_i(t)$ ,  $I_i(t)$ , and  $R_i(t)$  respectively. We denote the population of region  $i$  at time  $t$  as  $N_i(t) = S_i(t) + I_i(t) + R_i(t)$ . As time goes on, the number of individuals in these compartments can change. The epidemic dynamics are a function of both intra- and inter-region disease spread, described by the following differential equations involving a  $k \times k$  O-D matrix  $M$ , that are governed by three parameters:  $\beta$ ,  $\alpha$ , and  $\mu$ .

$$\begin{aligned}\frac{dS_i(t)}{dt} &= -\frac{\beta S_i(t)I_i(t)}{N_i(t)} - \frac{\alpha \beta S_i(t) \sum_{j \in [k]} \frac{M_{i,j}I_j(t)}{N_j(t)}}{N_i(t) + \sum_{j \in [k]} M_{i,j}} \\ \frac{dR_i(t)}{dt} &= \mu \frac{I_i(t)}{N_i(t)} \\ \frac{dI_i(t)}{dt} &= \frac{dS_i(t)}{dt} - \frac{dR_i(t)}{dt}\end{aligned}$$

The parameter  $\beta$  describes the frequency of mixing of intra-region populations in combination with the virality of the disease: how likely is an infected person to meet a non-infected person from the same region and transfer the pathogen? The  $\alpha$  term in the differential equations describes the frequency of mixing with inter-region visitors in relation to the frequency of mixing with intra-region residents: how likely is an infected person to meet a visitor and transfer the pathogen? The last parameter  $\mu$  describes the rate of recovery from disease: what proportion of the infected population recovers each day?

In principle,  $\beta$ ,  $\alpha$ , and  $\mu$  could all vary from region to region on the basis of a number of factors (for example, community cohesion, access to healthcare in the area, and mobility restrictions), but for simplicity we keep them constant across regions. To configure this model, we set  $\alpha = 1$ , meaning that there is no difference between the degree of mixing with residents and the degree of mixing with visitors. We additionally set  $\beta$  and  $\mu$  based on estimated dynamics from the COVID-19 pandemic. SIR estimates of  $\beta$  range between 0.06 and 0.39 in SIR models fit in 2020 in a set of nations; estimates of  $\mu$  range between 0.04 and 0.19 in the same set of studies<sup>68–70</sup>. For this case study, we set  $\beta = 0.10$  and  $\mu = 0.04$ .

The mobility-based SIR model relies additionally on mobility matrices  $M$  and a time-invariant regional population values  $N_1, \dots, N_k$ . To compare the intervention differences between the private and non-private regimes, we constructed non-private mobility matrices using the CDR from Afghanistan along with private versions with  $\epsilon \in \{0.1, 0.5, 1\}$  using our private O-D matrix algorithm at the trip level. We estimate the regional population values  $N_1, \dots, N_k$  by inferring a home location for each subscriber in the CDR dataset using the most common cell tower through which they placed calls between the hours of 8pm and 6am. These regional counts  $N_1, \dots, N_k$  could leak information about individuals, so in the differentially private version of the pandemic simulation we privatize the regional population counts  $N_1, \dots, N_k$  with the same value of  $\epsilon$  as the OD matrices.

As described in the results section, we use the mobility-based SIR model with daily origin-destination matrices to calculate epidemic curves (Fig. S1, Supplementary Information), and compare the accuracy of policy decisions taken when SIR models are deployed with private and non-private OD matrices.

### Geographic targeting of humanitarian aid after natural disasters and violent events

Our second simulated policy context is the targeting of humanitarian aid after natural disasters, violent events, and other shocks. As described in the results section, we calculate daily O-D matrices for time periods associated with two such shocks in our datasets: the Battle of Kunduz in Afghanistan in 2015, and the Lake Kivu Earthquake in Rwanda in 2008. We calculate daily O-D matrices for seven days following each event and calculate the two statistics using both private and non-private O-D matrices: (1) the total amount of out migration from the affected area, and (2) the top- $k$  districts with the most out migration. We assess the accuracy of both statistics using private data in comparison to non-private data. We also assess the sensitivity of these results to the choice of  $k$  (Fig. S3, Supplementary Information). To show the impact the utility of downstream policy decisions using individual-level privacy, we also perform geographic targeting simulations by selecting different values of  $T$  (Table S3, Supplementary Information).

## Data and code availability

The mobile phone datasets from Afghanistan and Rwanda contain detailed information on billions of mobile phone transactions. These data contain proprietary and confidential information belonging to a private telecommunications operator and cannot be publicly released. Upon reasonable request, we can provide information to accredited academic researchers about how to request the proprietary data from the telecommunications operator. Contact the corresponding author for any such requests.

Received: 15 June 2023; Accepted: 11 November 2024

Published online: 19 November 2024

## References

1. Wesolowski, A. et al. Impact of human mobility on the emergence of dengue epidemics in Pakistan. *Proc. Natl. Acad. Sci.* **112**, 11887–11892 (2015).
2. Wesolowski, A. et al. Quantifying the impact of human mobility on malaria. *Science* **338**, 267–270 (2012).
3. Bengtsson, L. et al. Using mobile phone data to predict the spatial spread of cholera. *Sci. Rep.* **5**, 1–5 (2015).
4. Lu, X., Bengtsson, L. & Holme, P. Predictability of population displacement after the 2010 Haiti earthquake. *Proc. Natl. Acad. Sci.* **109**, 11576–11581 (2012).
5. Pastor-Escuredo, D. et al. Flooding through the lens of mobile phone activity. *Proceedings of the 2014 IEEE Global Humanitarian Technology Conference (GHTC)* (2014).
6. Tai, X. H., Mehra, S. & Blumenstock, J. E. Mobile phone data reveal the effects of violence on internal displacement in Afghanistan. *Nature Human Behaviour* 1–11, <https://doi.org/10.1038/s41562-022-01336-4> (2022). Publisher: Nature Publishing Group.
7. Aiken, E., Bellue, S., Karlan, D., Udry, C. R. & Blumenstock, J. Machine Learning and Mobile Phone Data Can Improve the Targeting of Humanitarian Aid. *Nature* **603**. <https://www.nature.com/articles/s41586-022-04484-9> (2022)
8. Blumenstock, J. Machine learning can help get COVID-19 aid to those who need it most. *Nature* **581**. <https://doi.org/10.1038/d41586-020-01393-7> (2020).
9. Dobra, N., A. Williams & Eagle, N. Spatiotemporal detection of unusual human population behavior using mobile phone data. *PLoS ONE* **10** (2015).
10. Gundogdu, D., Incel, O., Salah, A. & Lepri, B. Countrywide arrhythmia: emergency event detection using mobile phone data. *EPJ Data Sci.* **5** (2016).
11. Kraemer, M. U. et al. The effect of human mobility and control measures on the COVID-19 epidemic in China. *Science* **368**, 493–497 (2020).
12. Warren, M. S. & Skillman, S. W. Mobility changes in response to COVID-19. *arXiv preprint. arXiv:2003.14228* (2020).
13. Chinazzi, M. et al. The effect of travel restrictions on the spread of the 2019 novel coronavirus (COVID-19) outbreak. *Science* **368**, 395–400 (2020).
14. Ilin, C. et al. Public mobility data enables COVID-19 forecasting and management at local and global scales. *Sci. Rep.* **11**, 1–11. <https://doi.org/10.1038/s41598-021-92892-8> (2021).
15. Lai, S. et al. Effect of non-pharmaceutical interventions to contain COVID-19 in China. *Nature* **585**, 410–413 (2020).
16. Oliver, N. et al. Mobile phone data for informing public health actions across the COVID-19 pandemic life cycle. *Sci. Adv.* (2020).
17. Thompson, S. A. & Warzel, C. Twelve million phones, one dataset, zero privacy. *The New York Times* **19**, 2019 (2019).
18. Ovide, S. The Nightmare of Our Snooping Phones. *The New York Times* (2021).
19. Boorstein, M., Iati, M. & Shin, A. Top U.S. Catholic Church official resigns after cellphone data used to track him on Grindr and to gay bars. *The Washington Post* (2021).
20. de Montjoye, Y.-A., Hidalgo, C., Verleysen, M. & Blondel, V. Unique in the crowd: The privacy bounds of human mobility. *Sci. Rep.* **3** (2013).
21. Privacy & Board, C. L. O. Report on the government's use of the call detail records program under the usa freedom act. *Privacy and Civil Liberties Oversight Board* (2020).
22. Hosein, G. & Nyst, C. Aiding surveillance: an exploration of how development and humanitarian aid initiatives are enabling surveillance in developing countries. Available at SSRN 2326229 (2013).
23. Taylor, L. No place to hide? the ethics and analytics of tracking mobility using mobile phone data. *Environ. Plann. D Soc. Sp.* **34** (2016).
24. de Montjoye, Y.-A. et al. On the privacy-conscious use of mobile phone data. *Sci. Data* **5**, 1–6 (2018).
25. Xu, F., Tu, Z., Zhang, P., Fu, X. & Jin, D. Trajectory recovery from ash: User privacy is not preserved in aggregated mobility data. In *Proceedings of the 26th World Wide Web Conference*, 1241–1250 (2017).
26. Abul, O., Bonchi, F. & Nanni, M. Never walk alone: Uncertainty for anonymity in moving objects databases. In *2008 IEEE 24th international conference on data engineering*, 376–385 (Ieee, 2008).
27. Primault, V., Mokhtar, S. B., Lauradoux, C. & Brunie, L. Time distortion anonymization for the publication of mobility data with high utility. In *2015 IEEE Trustcom/BigDataSE/ISPA*, vol. 1, 539–546 (IEEE, 2015).
28. Meyerowitz, J. & Roy Choudhury, R. Hiding stars with firewalls: location privacy through camouflage. In *Proceedings of the 15th annual international conference on Mobile computing and networking*, 345–356 (2009).
29. Shi, J., Zhang, R., Liu, Y. & Zhang, Y. Prisen: privacy-preserving data aggregation in people-centric urban sensing systems. In *2010 Proceedings IEEE INFOCOM*, 1–9 (IEEE, 2010).
30. Dwork, C., McSherry, F., Nissim, K. & Smith, A. Calibrating noise to sensitivity in private data analysis. In *Theory of cryptography conference*, 265–284 (Springer, 2006).
31. Mir, D. J., Isaacman, S., Cáceres, R., Martonosi, M. & Wright, R. N. DP-Where: Differentially private modeling of human mobility. In *2013 IEEE international conference on big data*, 580–588 (IEEE, 2013).
32. Qardaji, W., Yang, W. & Li, N. Differentially private grids for geospatial data. In *2013 IEEE 29th international conference on data engineering (ICDE)*, 757–768 (IEEE, 2013).
33. Aktay, A. et al. Google covid-19 community mobility reports: Anonymization process description (version 1.0). *arXiv preprint[SPACE]arXiv:2004.04145* (2020).
34. Machanavajjhala, A., Kifer, D., Abowd, J., Gehrke, J. & Vilhuber, L. Privacy: Theory meets practice on the map. In *2008 IEEE 24th international conference on data engineering*, 277–286 (IEEE, 2008).
35. Pratesi, F. et al. Privacy-aware distributed mobility data analytics. In *Proceedings of the 21th Italian Symposium on Advanced Database Systems, SEBD 2013* (2013).
36. Jiang, K., Shao, D., Bressan, S., Kister, T. & Tan, K.-L. Publishing trajectories with differential privacy guarantees. In *Proceedings of the 25th International conference on scientific and statistical database management*, 1–12 (2013).
37. Ghane, S., Kulik, L. & Ramamohanarao, K. Tgm: A generative mechanism for publishing trajectories with differential privacy. *IEEE Internet Things J.* **7**, 2611–2621 (2019).

38. Shao, D., Jiang, K., Kister, T., Bressan, S. & Tan, K.-L. Publishing trajectory with differential privacy: A priori vs. a posteriori sampling mechanisms. In *Database and Expert Systems Applications: 24th International Conference, DEXA 2013, Prague, Czech Republic, August 26–29, 2013. Proceedings, Part I 24*, 357–365 (Springer, 2013).
39. Li, M., Zhu, L., Zhang, Z. & Xu, R. Achieving differential privacy of trajectory data publishing in participatory sensing. *Inf. Sci.* **400**, 1–13 (2017).
40. Savi, M. K. et al. A standardised differential privacy framework for epidemiological modelling with mobile phone data. *medRxiv* 2023–03 (2023).
41. Shaham, S., Ghinita, G. & Shahabi, C. Differentially-Private Publication of Origin-Destination Matrices with Intermediate Stops. *arXiv preprint*. [arXiv:2202.12342](https://arxiv.org/abs/2202.12342) (2022).
42. Dwork, C., Kohli, N. & Mulligan, D. Differential privacy in practice: Expose your epsilons! *J. Privacy Confident.* **9** (2019).
43. Kermack, W. O. & McKendrick, A. G. A contribution to the mathematical theory of epidemics. *Proc. R. Soc. Lond. Ser. A* **115**, 700–721 (1927).
44. Goel, R. & Sharma, R. Mobility based sir model for pandemics-with case study of covid-19. In *2020 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*, 110–117 (IEEE, 2020).
45. Schroeder, A. The Kincade Fire: Population Movement and Displacement. *ArcGIS Story Maps* (2019).
46. CrisisReady. 7.8 and 7.5 Magnitude Earthquakes strike Turkey, Cause Mass Destruction and Growing Death Toll. *CrisisReady* (2023).
47. Afghan forces struggle as Taliban seeks northern stronghold. *Military Times* (2015).
48. Pyrgelis, A., Troncoso, C. & De Cristofaro, E. Knock knock, who's there? membership inference on aggregate location data. *arXiv preprint* [arXiv:1708.06145](https://arxiv.org/abs/1708.06145) (2017).
49. Naldi, M. & D'Acquisto, G. Differential privacy: An estimation theory-based method for choosing epsilon. *arXiv preprint*. [arXiv:1510.00917](https://arxiv.org/abs/1510.00917) (2015).
50. Hsu, J. et al. Differential privacy: An economic method for choosing epsilon. In *2014 IEEE 27th Computer Security Foundations Symposium*, 398–410 (IEEE, 2014).
51. Kohli, N. & Laskowski, P. Epsilon voting: Mechanism design for parameter selection in differential privacy. In *2018 IEEE Symposium on Privacy-Aware Computing (PAC)*, 19–30 (IEEE, 2018).
52. Ligett, K., Neel, S., Roth, A., Waggoner, B. & Wu, S. Z. Accuracy first: Selecting a differential privacy level for accuracy constrained erm. *Adv. Neural Inf. Process. Syst.* **30** (2017).
53. Whitehouse, J., Ramdas, A., Wu, S. Z. & Rogers, R. M. Brownian noise reduction: Maximizing privacy subject to accuracy constraints. *Adv. Neural Inf. Process. Syst.* **35**, 11217–11228 (2022).
54. Liu, J. & Talwar, K. Private selection from private candidates. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, 298–309 (2019).
55. Papernot, N. & Steinke, T. Hyperparameter tuning with renyi differential privacy. In *International Conference on Learning Representations* (2021).
56. Cummings, R. et al. Advancing differential privacy: Where we are now and future directions for real-world deployment. *Harvard Data Sci. Rev.* **6** (2024).
57. Kohli, N. & Laskowski, P. Differential privacy for black-box statistical analyses. *Proc. Privacy Enhancing Technol.* **3**, 418–431 (2023).
58. Zhu, K., Van Hentenryck, P. & Fioretto, F. Bias and variance of post-processing in differential privacy. In *Proceedings of the AAAI Conference on Artificial Intelligence* **35**, 11177–11184 (2021).
59. Kohli, N. *Leveraging differential privacy while attending to social and political commitments* (PhD diss. University of California, Berkeley, 2021).
60. Wang, Z. & Reiter, J. P. Post-processing differentially private counts to satisfy additive constraints. *Trans. Data Priv.* **14**, 65–77 (2021).
61. Pujol, D. et al. Fair decision making using privacy-protected data. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, 189–199 (2020).
62. Cohen, A., Duchin, M. & Matthews, J. *ϵ & Suwal, B* (Private numbers in public policy, Census, differential privacy, and redistricting. Harvard Data Science Review, 2022).
63. Zhu, K., Fioretto, F. & Van Hentenryck, P. Post-processing of differentially private data: A fairness perspective. In *International Joint Conference on Artificial Intelligence (IJCAI)* (2022).
64. *Suppression of Rates and Counts*. Centers for Disease Control and Prevention (CDC), United States Cancer Statistics (USCS) (2023).
65. *Dataquest Home. Data Suppression (Stability Reports) California Department of Education*.
66. Matthews, G. J., Harel, O. & Aseltine, R. H. Jr. A review of statistical disclosure control techniques employed by web-based data query systems. *J. Public Health Manag. Pract. JPHMP* **23**, e1 (2017).
67. *Suppression of Rates, Trends, and Counts*. U.S. An Ecosystem of Health Disparities and Minority Health Resources. Department of Health & Human Services National Institutes of Health.
68. Bertozzi, A. L., Franco, E., Mohler, G., Short, M. B. & Sledge, D. The challenges of modeling and forecasting the spread of covid-19. *Proc. Natl. Acad. Sci.* **117**, 16732–16738 (2020).
69. Bagal, D. K., Rath, A., Barua, A. & Patnaik, D. Estimating the parameters of susceptible-infected-recovered model of covid-19 cases in india during lockdown periods. *Chaos, Solitons & Fractals* **140**, 110154 (2020).
70. Lounis, M. & Bagal, D. K. Estimation of SIR model's parameters of COVID-19 in Algeria. *Bull. Natl. Res. Centre* **44**, 1–6 (2020).

## Acknowledgements

This work was supported by CITRIS and the Banatao Institute at the University of California, the Bill and Melinda Gates Foundation via the Center for Effective Global Action, DARPA and NIWC under contract N66001-15-C-4066, and the National Science Foundation under CAREER Grant IIS-1942702. Aiken further acknowledges funding from a Microsoft Research PhD Fellowship. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes not withstanding any copyright notation thereon. The views, opinions, and/or findings expressed are those of the author(s) and should not be interpreted as representing the official views or policies of the Department of Defense or the U.S. Government. We are grateful to Sveta Milusheva and Robert Marty for providing early feedback on this work, as well as Paul Laskowski and Deirdre Mulligan for helpful discussions.

## Author contributions

NK and JB conceived of the research idea. NK developed the private mobility matrix algorithm and proved each of the theorems in the paper. EA implemented the empirical measurements of accuracy and simulations of intervention accuracy. NK, EA, and JB wrote the manuscript.

## Declarations

### Competing interests

The authors declare no competing interests.

### Additional information

**Supplementary Information** The online version contains supplementary material available at <https://doi.org/10.1038/s41598-024-79561-2>.

**Correspondence** and requests for materials should be addressed to J.E.B.

**Reprints and permissions information** is available at [www.nature.com/reprints](http://www.nature.com/reprints).

**Publisher's note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2024