

# Measuring patterns of human behaviour through large-scale mobile phone data

Big Data for social sciences

Pål Sundsøy



Doctor Philosophiae

Faculty of Mathematics and Natural Sciences

Department of Informatics

UNIVERSITY OF OSLO

February 2017

© Pål Sundsøy, 2017

*Series of dissertations submitted to the  
Faculty of Mathematics and Natural Sciences, University of Oslo  
No. 1815*


ISSN 1501-7710

All rights reserved. No part of this publication may be  
reproduced or transmitted, in any form or by any means, without permission.

Cover: Hanne Baadsgaard Utigard.  
Print production: Reprosentralen, University of Oslo.

*‘The hope is that as you take the economic pulse in real time, you will be able to respond to anomalies more quickly’*

- Hal Varian, Google Chief Economist (Prof. Emeritus, University of California, Berkley)



Human migration patterns in Bangladesh, derived from mobility patterns in mobile phone datasets.

Visualisation by Pål Sundsøy.





# List of publications

## **1. Can mobile usage predict illiteracy in a developing country?**

Preprint available at *arXiv:1607.01337* [cs.AI]. 2016.

## **2. Deep learning applied to mobile phone data for Individual income classification**

Joint work with Bjelland, J., Reme B.A., Iqbal A. and Jahani, E.

Published in *International conference on Artificial Intelligence: Technologies and Applications* (ICAITA). Atlantic Press. 2016.

## **3. Mapping Poverty using mobile phone and satellite data**

Joint work with Steele, J.E., Pezzulo, C., Alegana, V., Bird, T., Blumenstock, J., Bjelland J., Engø-Monsen, K., de Montjoye, Y.A., Iqbal, A., Hadiuzzaman, K., Lu, X., Wetter, E., Tatem, A. and Bengtsson, L.

Published in *Journal of The Royal Society Interface* 14:20160690. 2017

## **4. The activation of core social networks in the wake of the 22 July Oslo bombing**

Joint work with Ling, R., Engø-Monsen, K., Bjelland, J. and Canright, G.

Published in *Social Networks Analysis and Mining ASONAM* (pp. 586-590). 2012.

## **5. Detecting climate adaptation with mobile network data: Anomalies in communication, mobility and consumption patterns during Cyclone Mahasen**

Joint work with Lu, X., Wrathall, D., Nadiruzzaman, M., Wetter, E., Iqbal, A., Qureshi, T., Tatem, A., Canright, G., Engø-Monsen, K. and Bengtsson, L.

Published in *Climatic Change*, 138(3-4), pp.505-519. 2016.

## **6. Comparing and visualizing the social spreading of products on a large-scale social network**

Joint work with Bjelland, J., Engø-Monsen, K., Canright, G. and Ling, R.

Published in *Influence on Technology on Social Network Analysis and Mining*, Tanel Ozyer et. al. Springer International Publishing. 2012.

## **7. Big Data-Driven Marketing: How Machine Learning outperforms marketers' gut-feeling**

Joint work with Bjelland, J., Iqbal, A., Pentland, A. and de Montjoye, Y.A.

Published in *International Conference on Social Computing, Behavioral-Cultural Modeling, and Prediction* (pp. 367-374). Springer International Publishing. 2014.

# List of publications not included in the thesis

## **8. Impact of human mobility on the emergence of dengue epidemics in Pakistan**

Joint work with Wesolowski, A., Qureshi, T., Boni, M.F., Johansson, M.A., Rasheed, S.B., Engø-Monsen, K. and Buckee, C.O.

Published in *Proceedings of the National Academy of Sciences*, 112(38):11887-92. 2015.

## **9. Improving official statistics in emerging markets using machine learning and mobile phone data**

Joint work with Jahani, E., Bengtsson, L., Bjelland, J., Pentland, A. and de Montjoye, Y.A.

In review, *EPJ Data Science*. 2017.

## **10. Unveiling Hidden Migration and Mobility Patterns in Climate Stressed Regions: A Longitudinal Study of Six Million Anonymous Mobile Phone Users in Bangladesh**

Joint work with Lu, X., Wrathall, D.J., Nadiruzzaman, M., Wetter, E., Iqbal, A., Qureshi, T., Tatem, A., Canright, G., Engø-Monsen, K. and Bengtsson, L.

Published in *Global Environmental Change* 38, pp.1-7. 2016.

## **11. Small and Even Smaller Circles: The Size of Mobile Phone-Based Core Social Networks in Scandinavia and South Asia**

Joint work with Ling, R., Canright, G., Bjelland, J. and Engø-Monsen, K.

Published in *Journal of Intercultural Communication Research* 41(3), pp.320-339. 2012.

## **12. Joy of Giving: Increasing Product Uptake by allowing customers to Forward**

Joint work with Bjelland, J., Canright, G., Iqbal, A., Grønnetvet, G., Norton, M. and Reme, B.A.

Article in preparation.

## **13. Handset-centric view of smartphone application use**

Joint work with Rana, J., Bjelland, J., Couronne, T., Wagner, D. and Rice, A.

Published in *Procedia Computer Science*, 34, pp.368-375. 2014.

## **14. Small circles: Mobile Telephony and the cultivation of the private**

Joint work with Ling, R., Bjelland, J. and Campbell, S.

Published in *The Information Society*, 30(4), pp.282-291. 2014.

## **15. The socio-demographics of texting: An analysis of traffic data**

Joint work with Ling, R. and Bertel, T.

Published in *New Media & Society*, 14(2), pp.281-298. 2012.

**16. Product adoption networks and their growth in a large mobile phone network**

Joint work with Bjelland, J., Canright, G., Engø-Monsen, K. and Ling, R.

Published in *Advances in Social Networks Analysis and Mining* (pp. 208-216). 2010

**17. A Social Network Study of Android VS Apple Smartphone battle**

Joint work with Bjelland, J., Ling, R., Engø-Monsen, K. and Canright, G.

Published in *Advances in Social Networks Analysis and Mining* (pp. 983-987). 2012

**18. Using Deep Learning to predict demographics from mobile phone metadata**

Joint work with Felbo, B., Lehmann, S., de Montjoye, Y.A. and Pentland, A.

Article in preparation.

**19. Diffusion of Information Through On-Demand Information Seeking Behavior**

Joint work with Riedl, C., Bjelland, J., Canright, G., Iqbal, A., Engø-Monsen, K., Qureshi, T. and Lazer, D.

In review, *PNAS*, 2017.

**20. Quantifying socio-economic segregation across continents using mobile phone metadata**

Joint work with Lind, J.T., Kotsadam, A., Reme, B.A. and Bjelland, J.

Article in preparation.

**21. Networks and income: Evidence from Individually Matched Income and Mobile Phone Metadata**

Joint work with Jahani, E., Saint-Jacques, G., Bjelland, J., Aral, S. and Pentland, A.

Article in preparation.

# Acknowledgement

This thesis is based on 7 (out of 21) selected papers related to the analysis of human behaviour using large mobile phone datasets. All publications were written after I started at Telenor Research 9 years ago. I am grateful for the opportunity given by my employer to do research on mobile phone datasets.

Doing research on ‘Big Data’ applied to social sciences has been my main focus during these years, and the interest in this subject has increased gradually throughout my time at Telenor.

Most of the publications are joint work with some very bright individuals. It has been a privilege for me to cooperate with all of you.

I would like to highlight several people. Yves-Alexandre de Montjoye, Eaman Jahani and Alex Pentland at MIT Media Lab; thanks for great research collaboration through several years. You have been a great source of inspiration. Linus Bengtsson and colleagues at Flowminder Foundation; I have really appreciated collaborating with you. Your passion and research dedication towards social good is admirable. Amy Wesolowski/Caroline Buckee at Harvard School of Public Health; thanks for involving me in the work on epidemic spreading. David Lazer at Northeastern; thanks for introducing us to all the great researchers in Boston. Guillaume Saint-Jacques and Sinan Aral at MIT Sloan; thanks for a good and interesting collaboration, bridging the gap between economics and Big Data science. Chris Riedl at Harvard/Northeastern; for a good collaboration on viral spreading. Of course, I also want to give a big thanks to my research colleagues at Telenor Research; Geoffrey Canright and Kenth Engø-Monsen, who introduced me to social network analysis. Rich Ling; you have softened up my (hard physics) view of the world with sociological insight. Johannes Bjelland; thanks for being such a great colleague for 9 years, including all the memorable trips around the world. Bjørn-Atle Reme and Gorm Grønnevet; thanks for the precious advice on economics and good discussions. Asif Iqbal; for helping us with data lobbying when everything seemed hopeless, and for your positivity. Thomas Couronne; thanks for all the nice talks around data visualisation. Juwel Rana and Jo Thori Lind; for valuable advice around thesis structure. I would also like to thank Xin Lu, Linnet Taylor, Sune Lehmann, Bjarke Felbo, Andreas Kotsadam, Taimur Qureshi, Gro Nilsen, Jessica Steele, Andy Tatem, Carla Pezzulo, Erik Wetter, David Wrathall, Weiqing Zhang, Mike Norton, Bjørn Hansen, Astrid Undheim, Christian Tronstad, Hanne-Stine Hallingby, Wenche Nag, Andrew Rice and Daniel Wagner for good discussions.

Oslo, Jan 2017



Pål Sundsøy

# Table of contents

<b>1</b>	<b>Introduction .....</b>	<b>1</b>
1.1	Authorship contributions .....	2
1.2	Research objective .....	2
1.3	Outline of the study .....	3
<b>2</b>	<b>Review of literature .....</b>	<b>5</b>
2.1	Big Data for social sciences.....	5
2.2	Conceptual framework .....	6
2.2.1	Mobile phone metadata – what are CDRs?.....	6
2.2.2	Human behavioural traces .....	7
<b>3</b>	<b>Methodology.....</b>	<b>13</b>
3.1	Research design.....	13
3.1.1	Data collection .....	13
3.1.2	Tools.....	14
3.1.3	Analytical framework.....	14
3.1.4	Social network analysis .....	15
3.1.5	Machine learning.....	17
<b>4</b>	<b>Key contributions .....</b>	<b>21</b>
4.1	Contributions to social good .....	21
4.1.1	Studies 1–3: Predicting socio-economic indicators .....	21
4.1.2	Papers 4–5: Understanding systemic shocks and disasters in society .....	27
4.2	Contributions to behavioural insight and marketing.....	31
<b>5</b>	<b>Discussion .....</b>	<b>36</b>
5.1	On the research question .....	36
5.2	Challenges and limitations .....	37
5.2.1	Privacy challenges.....	37
5.2.2	Analysis and interpreting data .....	38
5.2.3	The future of artificial intelligence in social sciences.....	40
<b>6</b>	<b>Conclusion.....</b>	<b>41</b>
	<b>Bibliography.....</b>	<b>42</b>
	<b>Publications.....</b>	<b>48</b>



# 1 Introduction

Every time a mobile phone customer makes a phone call, sends an SMS or generates Internet traffic there are traces left at the mobile operator. These digital traces, which have traditionally been used mainly for billing purposes, also have other benefits. This thesis is based on 7 publications where mobile phone logs are used to understand large-scale human behaviour. The applications range from improved economic and social wellbeing to marketing, and the level of granularity ranges from the individual to the level of society. Contributions include individual prediction of socio-economic indicators, poverty prediction and understanding human behavioural signals during disasters. Furthermore, how products spread over large social networks is investigated, and how this information can be exploited to perform large-scale marketing experiments. The size of the datasets analysed ranges from 500 million to 300 billion phone records.

Studies 1 and 2 develop scalable predictive models based on mobile phone logs to reliably infer the illiteracy status and income level of individuals. Such insight can be further aggregated to the geographical level to help vulnerable groups in society. Study 3 represents the first attempt to build predictive maps of poverty using a combination of mobile phone and satellite data, with Bangladesh as an example. Knowing where poor people live is a crucial component of poverty eradication, and this study complements expensive approaches that are entirely based on data from traditional surveys with low temporal frequencies. Studies 4 and 5 quantify people's behaviour during larger shocks in society: study 4 reveals human behavioural patterns through the eyes of mobile phone data during the 22<sup>nd</sup> July 2011 terror attack in Norway, while study 5 shows how people adapt to climate extremes by analysing financial, social and mobility behaviour from 5 million people during a cyclone event in Bangladesh. The aim is to gain understanding and to detect early-warning signals that can help prevent future disasters.

We know that social networks matter when purchase decisions are made. Study 6 is motivated by the question of how people adopt new products and services, and what role the underlying social network structure plays in the process.

Study 7 addresses how social network effects, together with discretionary income and timing, can be modelled and exploited in large-scale marketing experiments in Asia, targeting people with more personalised offers.

## 1.1 Authorship contributions

Authorship contributions for the included articles (I–VII) are specified below.

- I. **Can mobile usage predict illiteracy in a developing country?**  
PRS is the sole author.
- II. **Deep learning applied to mobile phone data for individual income classification**  
PRS took the initiative, performed the analysis and wrote the paper. JØB, BAR, AI and EJ commented on the draft and suggested improvements for the paper. PRS is the first author.
- III. **Mapping poverty using mobile phone and satellite data**  
JS held the main responsibility for the paper and the overall analysis. CP was responsible for survey data management, cleaning and processing, and interpretation and drafting of the final manuscript. PRS were responsible for management of the project from Telenor side, the CDR data management, cleaning, and production of CDR data, and interpretation and drafting of the final manuscript. JB, J.Bj, KE was responsible for interpretation, drafting, and production of the final manuscript. V.A., T.B., Y.M., X.L. and E.W. were responsible for interpretation and production of the final manuscript. A.I. and K.N.H. for handling of income survey data. AJT and LB were responsible for overall scientific management, interpretation and production of the final manuscript. All authors gave final approval for publication.
- IV. **The activation of core social networks in the wake of the 22nd July Oslo bombing**  
PRS took the initiative, collected the data, performed the analysis, interpreted the data and wrote the article. RL, JB, KEM, GC interpreted the data, revised the article and approved it. PRS is the first author.
- V. **Detecting climate adaptation with mobile network data: Anomalies in communication, mobility and consumption patterns during Cyclone Mahasen**  
XL held the main responsibility for the paper and overall analysis. DW positioned the paper from a development perspective. PRS collected and prepared the data, analysed the financial top-up data and revised the article. MN EW, AI, TQ, AT, GC, KEM, LB interpreted the data, suggested changes and gave the final approval of the article. PRS is the third author.
- VI. **Comparing and visualising the social spreading of products on a large-scale social network**  
PRS took the initiative, collected and visualised the data, performed the analysis and drafted the article. RL suggested changes and placed the paper in a sociological context. JB, KEM and GC interpreted, revised and approved the article. PRS is the first author.
- VII. **Big Data-driven marketing: How machine learning outperforms marketers' gut feeling**  
PRS and JB coded the experiment and made the model. Several trips to South-East Asia were required to prepare and run the experiment. AI arranged access to data sources. PRS performed the post-analysis and wrote the article. YA and AP suggested improvements, revised the article and positioned the paper. PRS is the first author.

## 1.2 Research objective

The key research question in this thesis can be formulated as follows:

**Question:** Apart from providing basic communication services, what kinds of positive impacts can we create for society and/or individuals using large-scale mobile phone datasets?



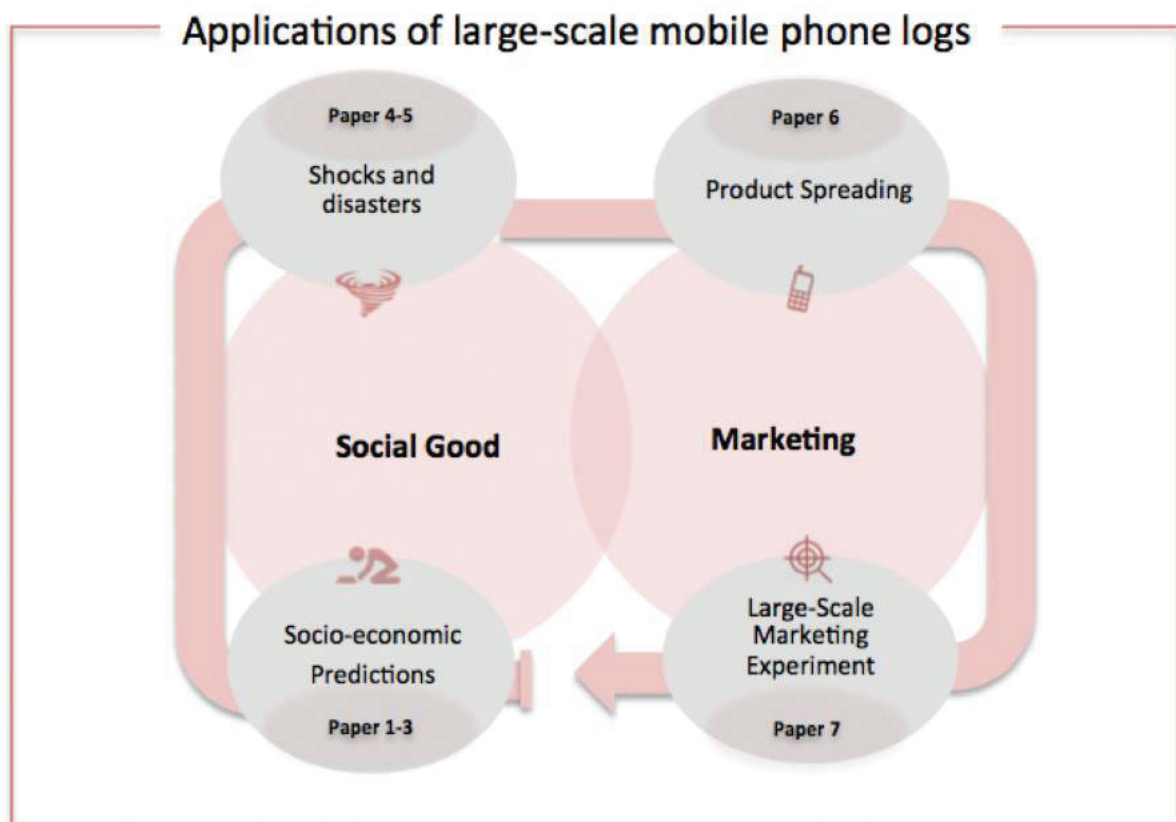
## 1.3 Outline of the study

The high-level dissertation roadmap is shown in Figure 1.1. This thesis argues that mobile phone data can be used to:

1. Inform socially beneficial policies.
2. Provide additional insights into human behaviour, with the aim of gaining:
  - I. A better understanding of human behaviour and interactions.
  - II. Better insights into human behaviour to improve marketing.

The publications in the first category include empirical studies that address challenges in society and how they can be tackled in a different way using mobile phone logs, while complementing existing approaches. The second category addresses how mobile phone logs can be used to obtain new behavioural insights and how such information can be used experimentally for marketing.

**Figure 1.1: Thesis theme overview. Publication number is given for each topic**



Chapter 2 reviews the literature, with a special emphasis on how mobile phone metadata has been used in the social sciences to date.

In Chapter 3 the research design, tools used and analytical framework are outlined.

Chapter 4 reviews the 7 papers ordered by topic, as shown in Figure 1.1, and introduces the problem, research findings and evaluation for each one.

Chapter 5 discusses the main research question, the limitations and challenges from a holistic point of view, followed by concluding remarks in Chapter 6. All publications are attached after the Bibliography.

## 2 Review of literature

The availability of large datasets, often referred to as ‘Big Data’, has opened the possibility of improving our understanding of society and human behaviour. The generation and use of large volumes of data is reshaping our social and economic landscapes, creating new industries, products and processes, and producing significant competitive advantages [1].

### 2.1 Big Data for social sciences

Recent research has found that countries could make much more use of data analytics in terms of economic and social benefits if governments did more to encourage investment in Big Data and to promote data sharing and reuse [2]. A consensual definition of Big Data is presented by De Mauro et al. : ‘*Big Data represents the information assets characterized by such a high volume, velocity and variety to require specific technology and analytical methods for its transformation into value*’ [3]. Naturally for Big Data, ‘size’ is a constantly moving target.

Studies have shown that Big Data has the potential to improve health policies [4], understand large-scale social networks [5], improve the efficiency of poverty prediction [6,7], run large-scale experiments and improve the understanding of urban development by, for instance, considering people’s mobility patterns [8]. Data-driven methods have also outperformed traditional marketing approaches [9]. By combining behavioural patterns in Big Data with traditional data it has also been shown to be useful in epidemic spreading predictions [10], and to compliment UN official statistics [11].

In summary, large-scale datasets of human behaviour have the potential to fundamentally transform performing social science research [12].

Traditionally, the broad range of social sciences have focused on explanatory models, whereas researchers from computational sciences have targeted predictive models or observational/‘found’ data [13]. Conventional statistical and econometric techniques such as regression often work well, but there are issues distinct to big datasets that may depend upon different approaches [14]. While large amounts of data will not overcome the selection problems that make causal inference difficult, this thesis shows that it can provide opportunities to gain understanding and run experiments on a scale that was previously

impossible in the social sciences. Recent literature also supports that there are many policy applications where causal inference is not central, or even necessary, and where machine learning is useful for solving prediction problems and generating high social impact [15]. One of the most promising and rich Big Data sources is mobile phone datasets. The mobile ecosystem is a major driver of economic progress and welfare globally. In 2014, the mobile industry generated 3.8% of global GDP [16]. Half of the world’s population now has a mobile subscription, and it has been reported that the global penetration rate in 2020 is expected to be 60%. De-identified mobile phone data is a promising source that has the potential to deliver near real-time information of human behaviour on both an individual and societal scale. The next chapter introduces the details of this promising source.

## 2.2 Conceptual framework

### 2.2.1 Mobile phone metadata – what are CDRs?

Whenever a mobile phone call or other transactions are made, a call detail record (CDR) is generated by the mobile operator [17]. A CDR contains, for instance, the start time and duration of a call, but does not provide any information about the content, and is therefore defined as mobile phone *metadata*. Other information recorded in CDRs is which cell tower the caller (and often recipient’s) phones were connected to at the time of the call. It is therefore possible to use the CDRs to approximate the location of a user. A sample of CDR fields is presented in Table 2.1. The location is usually found by coupling the cell ID to an external mapping table containing the actual positions (lon,lat) of the tower.

**Table 2.1: Example of set of CDR fields**

Calling party	Called party	Caller cell		Type	Call		
		ID	Call time		duration	IMSI	IMEI
91845206	92234065	6D45X	15.05.2016 14:24:50	Voice	200	4798X	5840X
91845206		A56DE	15.05.2016 20:10:13	Internet		4798X	5850X
92234065	91845206	A56DE	16.05.2016 15:40:25	SMS		4777X	6382X

Traditionally, CDRs have mostly been used for billing purposes and the maintenance of business. In recent years researchers from many cultures, data scientists, economists, social

scientists and public sector organisations have begun to explore additional applications of de-identified CDRs, where all personal information has been removed.

With regards to Telenor-owned mobile operators, where the data is taken from for use in this thesis, the number of *daily* stored CDR observations varies from 100 million to over 1 billion per country.

One additional promising source of metadata for social science research is airtime credit purchases, or so-called ‘top-up transactions’, which are used for recharging mobile accounts in pre-paid markets like Asia and Africa [18]. Each purchase contains the user ID, top-up amount, date/time of top-up and often the location information of the retailer used. See the example in Table 2.2. Retailers, used by customers as a touch-point to refill their account, can range from basic kiosks on the street level to stores in large shopping malls.

**Table 2.2: Example of a set of airtime purchase transaction fields**

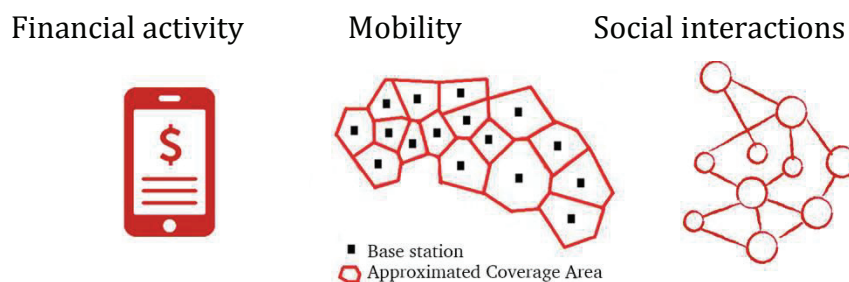
Buyer	Retailer	Retailer location	Time of purchase	Recharge amount
8849039482	880348403	6D45X	25.05.2016 20:10:13	300

Even if airtime purchases are stored separately they are often placed under the CDR umbrella, due to their transactional nature.

## 2.2.2 Human behavioural traces

Mobile phone metadata contains longitudinal digital traces of human behaviour which have proved valuable to guiding development policies and humanitarian action [19,20]. At least three dimensions can be measured: financial activity, mobility and social interactions.

**Figure 2.1: Dimensions measured by mobile phone metadata (CDR)**



**Financial activity:** When people in developing countries have more money to spend, they tend to spend a large portion of it on a luxury good such as cell phone communication, specifically by topping up their mobile airtime credit [21]. As we will see, several studies have already been run on socio-economic wellbeing using mobile phone data.

**Social interactions:** Social interactions include the number of ‘friends’ a person has, and how central the person is in the entire social network. In the context of mobile phone metadata the social network will be limited by the ‘call graph’ derived from phone logs, which is known to be a good proxy for the real social network [22]. The social interactions might also be used to understand how services and products diffuse on a large scale (or spread virally) throughout society [23].

**Mobility:** Since mobile phone users send and receive calls and messages through various cell towers, it is possible to reconstruct the movement patterns. This information may be used to understand daily rhythms of commuting to and from home, work, school and markets, but also have applications in modelling anything from a disease spreading to the movements of a disaster-affected population [24-26]. Typical derived personal-level features include most-used cell site, the radius of gyration (mobility radius) and total distance travelled within a given time period.

In addition to the three dimensions above, personal interests or content information, such as app usage, can be derived from on-board sniffer apps or deep packet inspection data [27]. Such information is however still very limited in use, and extends beyond the definition of metadata. There are also serious privacy implications [28].

The next section highlights examples of how mobile phone metadata has been applied to social sciences.

### *Socioeconomics*

Socioeconomics is the social science that studies how economic activity affects and is shaped by social processes. Eagle et al. quantified the correlation between network diversity and a population’s economic wellbeing [29]. The findings revealed that the diversity of individuals’ relationships is strongly correlated with the economic development of communities. The assumption that more diverse ties correlate with better access to social and economic opportunities was untested at the population level. It concludes that frequently making and receiving calls with contacts outside one’s immediate community is correlated with higher socio-economic class. Similar results were later verified by Jahani et al., which investigates

the differences between social networks of the rich and the poor based on individually matched income data [30].

Additionally, CDRs have also been shown to provide proxy indicators for assessing regional poverty levels, as proven by studies in Cote D'Ivoire [31] and Rwanda [6]. They can also complement national surveys in estimating the changes associated with a growing economy, by exploiting the relationships between socio-economic factors and cell-phone usage [32]. It has also been hypothesised that airtime purchases are correlated with socio-economic status, but it has been difficult to validate this with external reliable data [33]. Monitoring airtime purchases for trends can be useful for detecting early impacts of an economic crisis, as well as for measuring the impact of programmes designed to improve livelihoods and food security [34]. As shown later, someone's handset type might also be a good indicator of their economic wellbeing.

### *Population density*

Measuring population density in different regions can be explored by using the number of people who are calling each tower. By using CDRs, population density has been mapped out in France and Portugal [35] and Cote d'Ivoire [36]. In underdeveloped countries census data is costly and difficult to obtain, and existing data is often outdated, so CDRs can therefore provide updated information on the actual density of population in such regions. Population density has implications for economic growth and policies; the effects of population density and other socio-economic factors on poverty rates were studied in [37].

### *Unemployment*

The findings by Toole et al. highlight the potential of mobile phone metadata to improve forecasts of critical economic indicators, such as unemployment [38]. The researchers found that CDRs, specifically changes in mobility and social interactions, can be used to predict unemployment rates up to four months before the release of official reports and more accurately than using historical data alone. This research can potentially identify macroeconomic statistics faster and with much finer spatial granularity than traditional methods of tracking the economy.

### *Disasters and societal shocks*

Disasters initiate a complex chain of events that can disrupt the local economy, and in some cases the national economy [39]. It has been shown that CDRs can be used to monitor

extreme situations and predict the movements of people after natural disasters. Bagrow et al. studied the reactions of people to different emergency situations, such as a plane crash, bombing and earthquake [40]. During the Haiti earthquake, Bengtsson et al. identified mobility patterns by analysing CDRs, providing more post-analysis of population migration. They found that the destinations of people who left the capital during the first three weeks after the earthquake were correlated with the locations where they had strong social bonds [41,42]. This research indicated that relief efforts could be planned more precisely, as population movement patterns may be significantly more predictable than has been previously understood.

### *Infectious diseases*

Infectious diseases have serious consequences for people and the economies they affect. Diseases, such as dengue fever and cholera, often receive little attention because their real costs are poorly understood. Costs can be measured in loss of productivity, in dollars spent on healthcare interventions and in people's health and quality of life [43]. As CDR can help follow people's movements, these movements can also provide information about how a disease could spread across a country. The dynamics depend on the disease and how it can be transmitted, and therefore different models based on mobility have been suggested. Studies have validated the use of mobile phone data as a proxy for modelling epidemics [26]. Mobility patterns have been identified in Kenya, where regional travel patterns of millions of subscribers were mapped and related to areas in which malaria had a higher probability of spreading [10]. One shortcoming in this area of research comes from the current difficulty of gaining access to high-quality ground truth data to compare the results with. Gaining access to mobile phone data from more than one country has also proven to be very difficult, as exemplified by the Ebola outbreak [44]. The largest CDR analytics project on epidemics to date was performed on a dengue outbreak in Pakistan in 2013, where mobility patterns of 40 million subscribers were combined with climatic suitability indices and epidemiological data [45]. High-resolution fine-scale risk maps predicting epidemics were produced, providing a platform to allow local public health departments to prepare for epidemics.

### *Transportation and economic development*

Transportation has substantial effects on economic growth, but the relationship between transportation and the economy are poorly understood [46]. Several studies on CDRs using mobility patterns are based on the subject of transport planning, and avoiding traffic jams and



road accidents. Berlingiero et al. mapped new routes to decongest Abidjan's crowded roads, which could potentially reduce travel time by 10% [47]. Another study estimated the flow of residents between each pair of intersections in a city [48]. A recent study used CDRs to track commuters during peak morning rush hour in five cities [49]. People try to travel from home to work as quickly as possible, but simulations showed that as much as 30% of the total time lost to congestion is caused solely by what the authors refer to as selfish routing. They suggest that social route planning could make driving less problematic. For more information on the research in this field the reader is referred to [50].

## **Marketing**

Marketing is a wide term that touches upon many disciplines. From a societal point of view, marketing provides the link between the material requirements in society and its economic response patterns. One factor that influences marketing strategies is social forces and peer-influence [51]. As social networks are becoming more explicit through other technologies, understanding how peer-influence creates and sustains behavioural congestion is also becoming more feasible [52]. For instance, Trusov et al. has demonstrated that, on average, approximately one fifth of a user's friends actually influence his or her activity level on a specific website [53].

In particular, mobile phone metadata provides the opportunity to quantify the effects of social interactions on marketing (and vice versa) on a scale that has never been done before. It is known that the neighbourhood of an individual influences their decisions [54]. A study by Hill et al. quantified the neighbours' probability of adopting an undisclosed technological service [55], and Ahorony et al. showed that common app installations were overrepresented for pairs of users who often have physical meetings [56]. Furthermore, Risselada et al. showed that the neighbours' influence on product adoption evolves over time, depending on the time since introduction of the product in the market [57].

Product adoption has also been recently studied in other contexts, such as the adoption of Mobile Money [58]. Billions of people around the world live without access to banks or other formal financial institutions, and mobile money platforms, which deliver basic financial services over the mobile phone network, are believed to improve the lives of the poor. Unfortunately in many countries the adoption rates are still very low. A recent research study addresses how machine learning can be used to predict passive and active mobile money usage, using behavioural information from CDRs [59]. The results highlight key correlations of mobile money use in three development countries, as well as the potential for such

methods to drive adoption. However, the models developed in one country do not perform very well in other countries, which may indicate, in the context of mobile money, that each population has a unique signature in terms of which metrics are good predictors of adoption.

## 3 Methodology

This dissertation combines methods from multiple disciplines, including social network analysis, statistical analysis, machine learning and visualisation. The approach is data-centric, where inductive reasoning is used to make generalisations from specific observations.

Extensive travelling, especially in South-East Asia, has been required to collect and analyse the datasets. This chapter elaborates on the data collection, tools and analytical framework used.

### 3.1 Research design

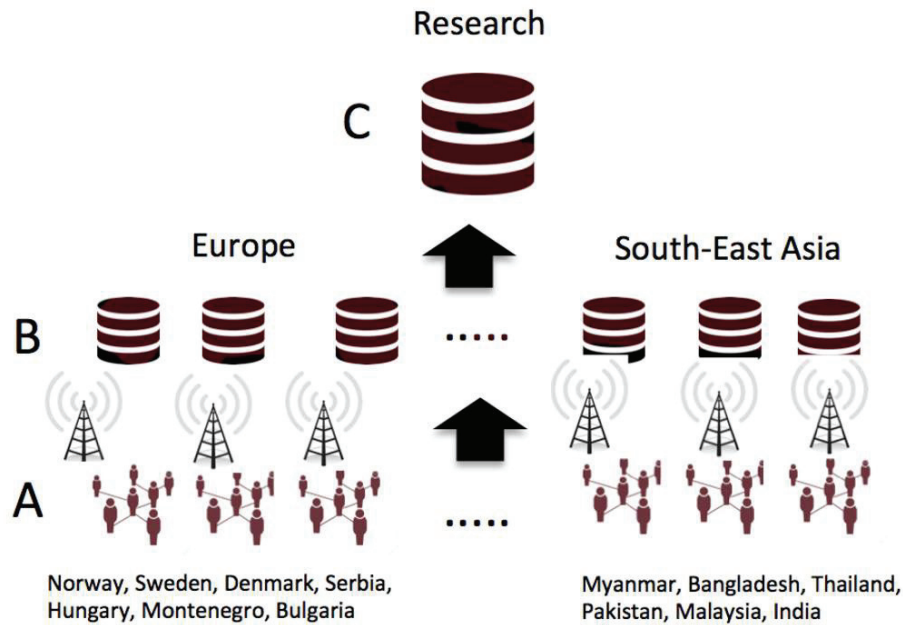
#### 3.1.1 Data collection

The data was collected from subsidiaries of the Telenor Group [60]. Telenor Group has operations in 13 countries in South-East Asia, Eastern Europe and Scandinavia (Fig 3.1), covering over 200 million mobile subscribers. Local data warehouses store the raw data, which is maintained by business intelligence teams. Research datasets are de-identified and either analysed locally, or transferred to a research data warehouse in Norway. The collection process is shown in Figure 3.2. In the case of external research collaboration, de-identified data are shared under special non-disclosure agreements.

**Figure 3.1: Global presence of Telenor in 13 countries (blue)**



**Figure 3.2: Data collection.** Network information is collected from mobile users (A) and stored in local data warehouses (B). De-identified information (CDR, purchases) is transferred to the research data warehouse via a secure channel



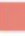



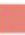




















### 3.1.2 Tools

A combination of several tools is used in this thesis. All analysis originates from second-resolution raw data, where processing requires a significant amount of computational power. Most datasets are pre-processed in-database using SQL, either on Oracle, Teradata or PostGres platforms, on high-performance Linux servers [61]. Considering the size of the datasets, SQL is also used for efficiency purposes for specific analytical and algorithmic tasks. Other tools include Python, R and specific tools for data analysis and machine learning [62,63]. For visualisation purposes, tools include Autodesk Maya for 3D modelling [64], and Gephi/Cytoscape for network visualisations [65]. QGIS is mainly used for spatial analysis and geographic visualisations [66].









### 3.1.3 Analytical framework

The analytical framework is considered as multi-disciplinary. Fig 3.3 provides an overview of the type of data sources and the *main* methodology used in each of the included publications.

**Figure 3.3: Per-paper overview of data sources and methodology**

Topic	Study	Data sources used	Main Methodology
Socio-economic prediction	1	  	
	2	  	
	3	   	
Shocks and disasters	4		
	5	  	
Product Spreading	6		
Marketing Experiment	7	  	

Data sources	Method
 Call Detail Records (CDR)	 Supervised Machine Learning & Prediction
 Financial airtime purchase data (top-up)	 Social Network Analysis
 Survey data	 General Statistical Analysis
 Subscription data (e.g. internet packages)	
 Spatial layers (e.g. satellite data, weather information)	

Papers 1, 2, 3 and 7 use supervised machine learning and/or prediction methods to infer a function from the training data that is further used for prediction. These studies make use of surveys or other subscriber information as ground truth information. An important note is that the feature generation processes (which generate the variables) are often based on other types of analysis, including mobility and social network analysis. Paper 3 uses techniques from machine learning, such as cross-validation and test-set, but within a Bayesian statistical modelling framework [67]. The methodology in papers 4 and 6 are mainly social network analysis. Paper 5 applies general statistical analysis and anomaly detection. All of the studies use CDR data as the main data source, while some of the studies make use of additional data sources such as financial airtime purchase information, subscription information and spatial layers, including satellite information. Data visualisation has also been vital, but is not mentioned in the methodology as shown in Figure 3.3.

### 3.1.4 Social network analysis


Mobile operators often have access to a huge portion of the social network in a given country, and these sources are therefore good candidates for social science research, e.g. for studying

social influence and purchase decisions. Understanding the nature of relationships and connections between entities is important for understanding a range of phenomena throughout multiple disciplines. Social network analysis (SNA) has broad and successful applications in economics epidemiology, sociology, biology and criminology [68]. The building block in the field of SNA is graphs, residing in graph theory, employed to represent the structure of interactions among people or any type of entities [69]. From a network perspective, it is the structure of the network and how the structural properties affect behaviour that is informative, and not simply the characteristics of the actors in the network. Analysis of large social networks is a non-trivial task that introduces challenges due to long processing time and large computational resources.

### *Basic elements in network analysis*

A social network is defined as graphs representing social relationships between people or organisations. Each node, also called a vertex or an actor, in a graph represents an individual person or a group of people. The connection between two individual nodes is referred to as an edge or tie. Two important concepts are components and centrality.

**Figure 3.4: Example social network**



**Social component:** A component is a portion of the network where all actors are connected, directly or indirectly, by at least one tie. By definition, each isolate is a separate component. Figure 3.4 consists of a social network of three social components. Node G is situated in an isolated component, while component 3 is the largest connected component (LCC) in the network.

**Centrality:** Centrality measures identify the most prominent actors, that is those who are extensively involved in relationships with other actors. The ‘importance’ of actors in networks is indicated by centrality. The most used centrality measures include degree

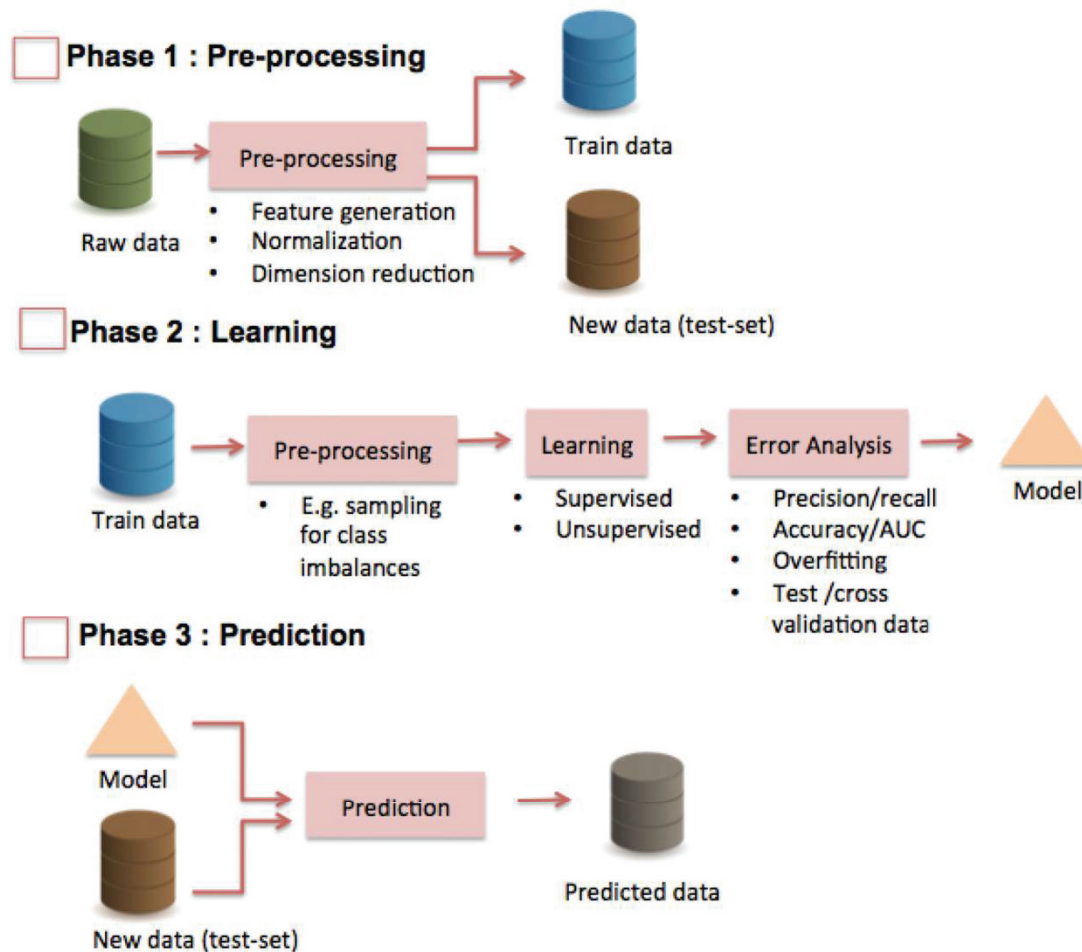
centrality, closeness centrality, betweenness centrality and eigenvector centrality [70]. By using Figure 3.4 as an example, we see that Nodes H and A both have three close friends each, indicating a degree centrality of three. If we use another centrality metric, such as eigenvector centrality (EVC), we will see that Node H suddenly has a higher EVC than A. The reason is that EVC takes into account the global network, where the friends of friends also matters.

From mobile phone datasets we can study a weighted social network, where the weight (or intensity)  $w_{ij}$  of an edge connecting person  $i$  and person  $j$  is defined as the aggregated time that the two users spend talking to one another. Most often, in communication networks the edge weight ( $i,j$ ) is taken as the total number of calls, or as the aggregated duration of calls between  $i$  and  $j$  during the period under investigation. Previous studies have shown that they give an equivalent quantification of edge weight [71].

### **3.1.5 Machine learning**

Machine learning picks its algorithms from different academic disciplines [72]. It is closely related to and often overlaps with computational statistics, which is a discipline that also focuses on prediction-making through the use of computers, and generalising from examples [73]. It is also considered as a sub-field of artificial intelligence. Practically, machine learning can be described as the algorithmic part of a data mining process, where the data preparation step is often the most tedious one. This type of data mining process model, which describes commonly used approaches to tackle problems, can be best visualised with the CRISP-DM framework [74]. A more detailed scheme of the isolated machine learning process is visualised in Figure 3.5. Here the process has been divided into 3 phases. The first phase is the pre-processing phase, where the features are generated from raw data, and the data are split into train and test.

**Figure 3.5: Process of machine learning**



The second phase is the computational learning phase, where techniques such as cross-validation are used to generate the final model [75]. Cross-validation is a technique for assessing how the results will generalise to an independent dataset, prevents the model from overfitting and provides an insight on how the model will generalise to an independent dataset, which is used in phase 3. There are also several ways to evaluate the performance of a machine learning model: popular choices include accuracy, recall, precision and ‘area under the curve’ (AUC) [76]. For instance, accuracy measures the proportion of true positives and negatives in the whole dataset. It is calculated for a given threshold; for example, logistic regressions return positive or negative depending on whether the logistic function is greater or smaller than a threshold, usually set to 0.5 by default. AUC measures how true positive rate (recall) and false positive rate trade off, and is an evaluation of the classifier as threshold varies over all possible values. The interpretation of AUC is the probability that a randomly chosen positive example is ranked above a randomly chosen negative example, according to



the classifier's internal value for the examples. The third phase is where the model is applied to unseen data (test-set), predictions are made and performance metrics are evaluated. If the performance is significantly lower than the train set, an overfitted model is used, and a re-calibration of the model is needed.

By the learning style (phase 2), the machine learning algorithms can be mainly divided into the following types: supervised learning and unsupervised learning.

### **Supervised learning**

In supervised learning, each sample in the dataset is a pair of an input vector and an external output vector (or value) that we are trying to predict [77]. Supervised learning is closely related to regression or classification in econometrics. By analysing a training set under a supervised learning algorithm an inferred function can be generated. The inferred function, e.g. the training model, can be then used to map or predict new samples. Both classification and regression are supervised learning where there is an input vector  $X$  and an external output  $Y$ , and the task  $T$  is to learn the experience  $E$  from the input  $X$  to the output  $Y$ . Typical supervised learning algorithm types are shown in Table 3.1.

**Table 3.1: Examples of supervised learning algorithms [78-81].**

Linear regression	Non-linear regression and classification	Regression and classification trees
<ul style="list-style-type: none"> <li>• Ordinary linear regression</li> <li>• Partial least squares regression</li> <li>• Penalised regression</li> </ul>	<ul style="list-style-type: none"> <li>• Multivariate adaptive regression splines</li> <li>• Support vector machines (SVM)</li> <li>• Artificial neural networks (incl. deep learning)</li> <li>• K-nearest neighbours</li> </ul>	<ul style="list-style-type: none"> <li>• Bagging tree</li> <li>• Random forest (RF)</li> <li>• Boosted trees/gradient boosted trees</li> </ul>

### **Unsupervised learning**

In unsupervised learning there is no external output (or label), and we only possess the input vector. The aim is to find similarities among samples in the unlabelled dataset. Typical algorithms include clustering (e.g. k-means), latent variable models and blind signal separation techniques [82]. Principal component analysis (PCA) is an example of a blind

signal separation technique that can be used to explain key features of the data, and reduce the data from a high-dimensional space. Often the feature selection itself is embedded into learning algorithms. Unsupervised learning is closely related to density estimation and clustering in econometrics. In this thesis supervised machine learning is mainly considered, as access to labelled data is available, and it is therefore possible to evaluate predictions.

### *A comment on econometrics vs. supervised machine learning*

The main objective in supervised machine learning is to provide accurate *predictions* of the variables of interest. Even though these techniques are extremely powerful for forecasting, it can be very difficult to interpret the underlying structure implied by them.

On the contrary, the main objective in econometrics is to provide an *explanation* of various observed outcomes. The goal is often to produce reliable estimates of parameters that describe economic systems, to provide an understanding of the underlying process that determines equilibrium outcomes. The estimation process is based on conditions implied by economic theory. Such a structural approach is beneficial when we want to know what happens when ‘the world’ changes (used for e.g. auctions, pricing). Machine learning can not easily predict the effect of intervention (how  $y$  changes as some  $x$  change).

Hal Varian argues that there are several things that econometricians can learn from machine learners and vice versa [14]. Machine learning introduces the train-test-validate (including cross-validation) concept to avoid model overfitting, and there are several non-linear estimation techniques, as mentioned above. Variable selection methods are also developed to deal with large amounts of data.

However, casual inference including confounding and instrumental variables is seldom considered in machine learning, even though some efforts have been made to combine the two approaches [83,84]. In addition, time series are often decomposed into trend and seasonal components to look at deviations from expected behaviour. The concept of cross-validation does not work directly on time series. Kleinberg et al. show how machine learning adds value over traditional regression approaches in solving prediction problems, and argue that causality is not always important when dealing with policy applications [15].

## 4 Key contributions

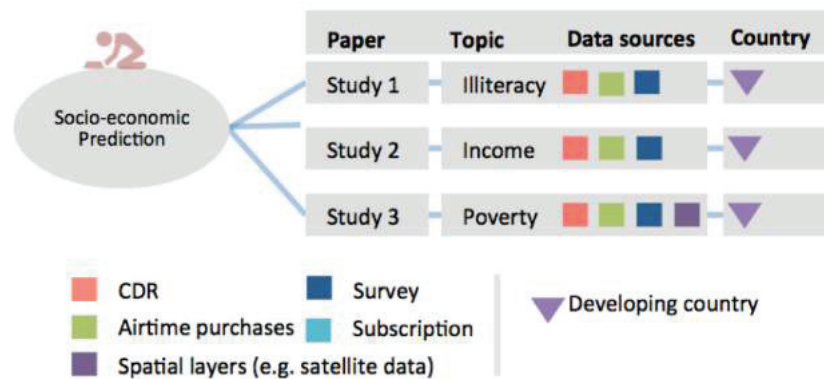
This chapter summarises the main findings in each of the 7 papers included in this thesis. Section 4.1 focuses on using mobile phone logs in a context that can benefit society as a whole. Contributions include individual prediction of socio-economic indicators, poverty prediction and understanding human behavioural signals during disasters. Section 4.2 investigates how products spread over large social networks, and further how to use these key findings in a large-scale marketing experiment in Asia, giving people access to more personalised offers.

### 4.1 Contributions to social good

#### 4.1.1 Studies 1–3: Predicting socio-economic indicators

Section 2.2.2 discussed how mobile phone metadata has been used to study socio-economic behaviour. Here we build on these findings; in study 1 and 2 prediction models are developed that can be used to infer individual characteristics of users, such as illiteracy status and income, by looking only at the users' mobile phone behaviour. The fact that most phones in the developing world are prepaid means that the data lacks very basic information about the individual. This prevents numerous uses of this data in development economics research and social sciences. More importantly, it prevents the development of humanitarian applications such as the use of mobile phone data to target aid towards the most vulnerable groups during a crisis. For development purposes insight can be learned on the individual scale, before it is aggregated on a spatial scale of mobile towers.

**Figure 4.1: Studies 1–3: topic, data sources and country**



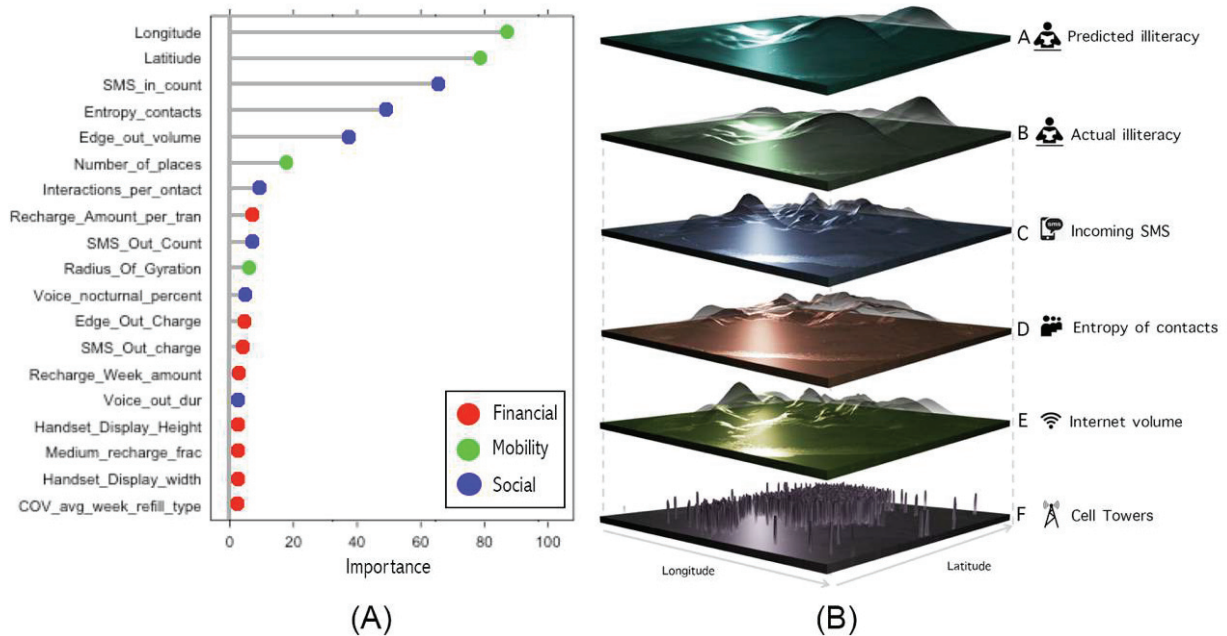
Study 3 addresses how mobile phone data can be used to predict multiple dimensions of poverty in Bangladesh using both mobile phone metadata and additional spatial information, such as satellite information, as model input. As indicated in Figure 4.1, all the studies use CDR and survey information as their main data sources.

#### *Study 1: Can mobile usage predict illiteracy in a developing country?*

The aim with study 1 is to investigate whether large-scale mobile phone metadata, in the form of CDRs and airtime purchases (top-up) can support quantifying individual and spatial illiteracy in a developing country. Geographical mapping of illiteracy is crucial to know where illiterate people are, and where to distribute resources to improve education. In underdeveloped countries such mapping can be based on out-dated household surveys with low spatial and temporal resolution. One in five people worldwide struggle with illiteracy, and it is estimated that illiteracy costs the global economy more than \$1 trillion dollars each year [85].

By deriving a broad set of mobile phone indicators reflecting users' financial, social and mobility patterns, as introduced in section 2.2.2, we show how supervised machine learning can be used to predict individual illiteracy. On average the model performs 10 times better than random guessing with 70% accuracy. Feature investigation indicates that the most frequent cell tower and incoming SMS are the superior predictors, followed by diversity of communication partners and Internet volume (Figure 4.2a). Furthermore, the investigation extends to how individual illiteracy can be aggregated and mapped geographically at cell tower resolution (Figure 4.2b).

**Figure 4.2: a) Most important predictors from mobile phone metadata data for predicting illiteracy; b) geographical mapping of illiteracy, top predictors and the cell tower distribution in one major city**



This study shows how illiteracy can be predicted from mobile phone logs, purely by investigating users' metadata. An important policy application of this work is the prediction of regional and individual illiteracy rates in underdeveloped countries where official statistics are limited or non-existing. Further work is required to investigate the findings up against population illiteracy and also verify the robustness of indicators in other countries.

### ***Study 2: Deep learning applied to mobile phone data for individual income classification***


One promising supervised classification method for social sciences is deep learning, sometimes referred to as 'deep neural networks', as introduced in section 3.1.5. For social sciences the main advantage of deep learning is to avoid cumbersome feature engineering, and let the algorithm itself decide the features (based on near raw data as input).

Deep learning has had breakthroughs in computer vision [86] and speech recognition [87]. The aim of study 2 is to understand whether deep learning can be beneficial for useful prediction tasks on mobile phone data, where classic machine learning algorithms are often under-utilised due to time-consuming country and domain-specific feature engineering (as shown in paper 1). Specifically the aim is to see how well the socio-economic status of an

individual can be predicted, with a comparison to traditional data mining models as a benchmark.

A simple deep learning architecture is implemented and is compared with traditional data mining models as benchmarks. On average this model achieves 77% AUC on test data using location traces as the sole input. In contrast, the benchmarked state-of-the-art data mining models include various feature categories such as basic phone usage, top-up pattern, handset type, social network structure and individual mobility. The traditional machine learning models achieve 72% AUC in the best-case scenario. Figure 4.3a shows the predictive performance, as measured by the AUC on test set, where the true positive vs. false positive rate is plotted for deep learning (DL), gradient boosting machines (GBM) and random forest (RF). The top predictive indicators in the random forest model and given in 4.3b and coloured by their respective variable family.

**Figure 4.3: a) Income prediction; predictive deep learning (DL), gradient boosting machines (GBM) and random forest (RF); b) top predictors for random forest model**



Even though the traditional model is not meant to be interpretable, Figure 4.4b gives some insight into the most predictive features. We especially note the importance of location dynamics, handset brand and airtime credit purchase patterns for predicting income.

*Location dynamics:* Where the user spends most of his time is a good signal of his income. This indicates that our models have detected regions of low economic development status. This is also in line with the deep learning model.

*Handset brand:* In the country of the study, minimal and more affordable handset brands are very popular among the lower income quantiles, while expensive smartphones are considered as a huge status symbol.

*Top-up pattern:* Interestingly, the recharge amount per transaction is more predictive than the total recharge amount. It can be observed that individuals from the lower income quantiles usually top-up with lower amounts when they first fill up their account.

The approach suggests that deep learning approaches could be an effective tool for predicting economic indicators based on mobile communication patterns. The disadvantage is that this is a more ‘black box’ approach and harder to interpret than traditional models.

### ***Study 3: Predicting multiple dimensions of poverty using mobile phone and satellite data***


In 2015, approximately 700 million people lived in extreme poverty, defined as living on less than \$1.90 a day [88]. To end poverty in all its forms everywhere is one of the selected targets of the UN Sustainable Development Goals [89]. As mentioned in section 2.2.2, CDRs have been shown to provide proxy indicators for assessing regional poverty levels in Cote D’Ivoire and Rwanda. Paper 3 builds on these approaches, and extends the approach by also including spatial layers such as satellite information, in addition to predicting multiple dimensions of poverty, with Bangladesh as an example.

Eradication of poverty requires national and subnational quantification and monitoring over the next 15 years, and the challenge is to establish appropriate, effective and timely measurements. Existing approaches to estimate multi-dimensional poverty rely on census data collected with limited temporal frequency. Alternate measures are needed to update estimates in the time between censuses. Here the aim is to investigate whether mobile phone data, combined with satellite data, can complement existing approaches for predicting multiple dimensions of poverty. An objective is to know which sources are most promising in rural vs. urban areas.

In this work, a Bayesian geostatistical modelling framework is used, combining data from GPS-located household surveys, satellite and other spatial layers, and mobile phone metadata to predict asset-, consumption- and income-based metrics of poverty at high resolution. This reveals that models employing a combination of mobile operator data and satellite variables provide the highest predictive power and lowest uncertainty (highest for wealth index:  $r^2=0.78$ ).



**Figure 4.4: National level prediction maps for mean wealth index. Maps were generated using mobile phone features, remote sensing data and Bayesian geostatistical models. The lower map restricts the focus on the poorest cells.**



National, urban and rural models were built to predict poverty rates from three separate geo-referenced household survey datasets. Indicators such as night-time lights, transport time to the closest urban settlement and elevation were important, both nationally and in rural models; climate variables were also important in rural areas. Distance to roads and waterways were significant within urban and rural strata. In general, the addition of CDR data did not change the selection of satellite features at any level. Top-up features derived from recharge amounts and tower averages were significant in every model. Percentage of nocturnal calls, and count and duration of SMS traffic were significant nationally. Mobility and social network features were important at all strata, but only in rural models when combined with satellite data. In urban areas, SMS traffic was important, whereas multimedia messaging and video attributes were key in rural areas. The outputs correlate strongly with previous poverty estimates for Bangladesh, highlighting the value of such cell phone-satellite-driven models in producing high-resolution poverty maps that can be rapidly updated. The findings based on

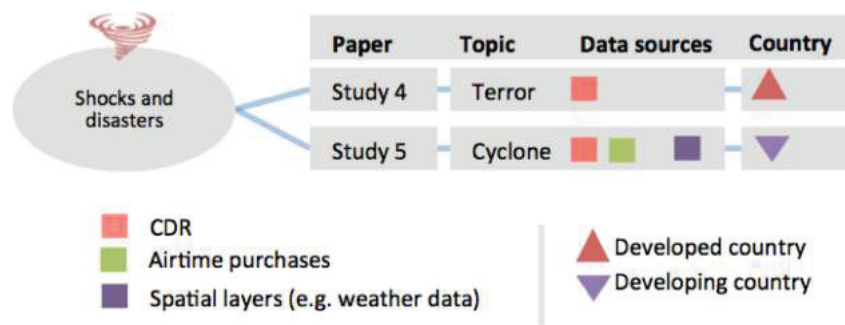


this research can be utilised for real-time monitoring and decision-making to more effectively reduce poverty.

#### 4.1.2 Papers 4–5: Understanding systemic shocks and disasters in society

A shock is an unexpected or unpredictable event that affects the economy, either positively or negatively. Shocks are typically produced when accidents or disasters appear. Systemic shocks will prompt hundreds or thousands of individuals or households to react in roughly similar ways. Section 2.2.2 discussed how mobile phone data have been used to study shocks and disasters, such as in the case of Haiti [41].

**Figure 4.5: Studies 4–5: topic, data sources and country**



The studies introduced in this section are related to two extreme events, as shown in Figure 4.5. The first study uses CDR information to investigate the terror incident that hit Oslo, Norway, in 2011. The next study investigates an extreme weather event – cyclone Mahasen, which hit Bangladesh in 2013.

##### *Study 4: The activation of core social networks in the wake of the 22nd July Oslo bombing*


Study 4 examines human behavioural patterns on Friday 22nd July 2011 when a powerful bomb exploded in Regjeringskvartalet (the centre of national administration) in central Oslo, Norway. It killed 8 people and seriously injured almost 100 others.

Earlier work has presented qualitative results on the 9/11 catastrophe, emphasising the need to reach out to the closest tie [90]. Another study examined the geographical distribution of traffic after a terrorist bombing in Israel [91]. This work differs from the others in the sense that large-scale mobile phone logs are used to understand communication patterns around the

disaster. In such events we feel the need to check the wellbeing of family and friends, to organise assistance, and to make sense of the situation. Mobile phone logs can illuminate the ways these needs are met in the face of disaster.

Empirical mobile traffic data illuminate exceptional behaviour immediately after the bombing in Oslo; in the minutes after the bombing people called ties that were close socially and perceived to be in danger, that is, people who were close to the bombing point. The main findings: (1) individuals first focus on their single closest contact ('best friend'), but soon after switch to spending more mobile communication resources than average on contacts ranked 2–5; (2) a large increase (over typical) in traffic is clear to and from, and not least within, the affected area (Oslo). In some cases this was more than a 300% increase immediately after the bombing. Interestingly, a marked increase in traffic also occurred for relationships where both persons were outside Oslo. All of these results illustrate the importance of social contact in this highly unusual situation.

**Figure 4.6: Fraction of active subscribers contacting their 1–5 closest relations vs. time for all active subscribers. All curves are normalised by average communication for the same contact number (1–5)**



This paper underlines how the mobile phone is an instrument of the intimate sphere. The situation on the 22nd July in Oslo is a prime example of an unexpected situation where

individuals use the mobile to get critical information on their loved ones in their core network. The study has several limitations. These include the fact that the data are taken only from one operator in Norway, and the data is therefore not generalisable to all of Norway. The location-based data use also the postal code of the subscriber, and not the actual location of the phone at the time of the bombing. This is a result of privacy issues.

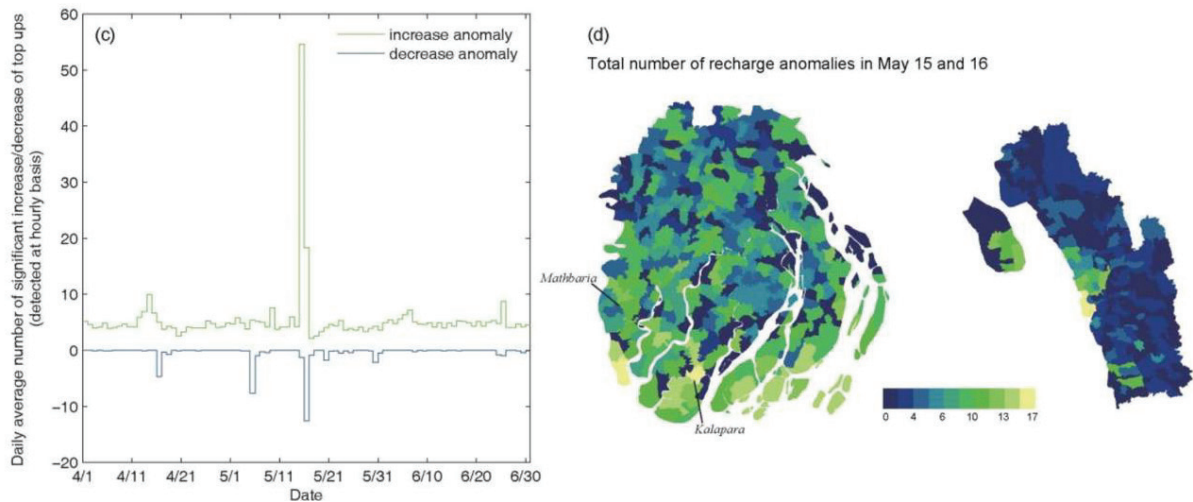
***Study 5: Detecting climate adaptation with mobile network data: anomalies in communication, mobility and consumption patterns during Cyclone Mahasen***

Other shocks in society include natural disasters. Extreme weather events have always had and will continue to have significant consequences for society and the economy. Climate projections indicate that changing extreme weather patterns are very likely to increase exposure to those events. Researchers have also found that a temperature change will leave the average income around the world 23% lower in the year 2100 than without climate change [92]. The following paper quantifies the impact of an extreme weather event in Bangladesh using CDR data. Weather events in Bangladesh already have a major impact on the economic performance and livelihoods of millions of poor people, and climate change is likely to drive migration from environmentally stressed areas.

In this study mobility, economic and social patterns are analysed during cyclone Mahasen, which hit Bangladesh in 2013. The aim is to investigate whether data from mobile phone metadata may be a useful tool to prioritise locations in which rapid needs assessment is performed after a cyclone. The aim is also to investigate whether anomaly detection can help to understand ‘signals’ of response in the population exposed to the cyclone.

The results show that anomalous patterns of calling frequency correlate with rainfall intensity at the local scale, likely providing a spatiotemporal indicator of users’ physical exposure to the storm.

**Figure 4.7: The temporal and spatial distribution of anomalies in airtime purchase anomalies. The threshold detection was set at three standard deviations from the mean of baseline**



In addition, the results show that mobile recharge purchases increase in impact zones before landfall, representing preparations for potential environmental threats. The temporal and spatial recharge purchase anomalies are shown in Figure 4.7. Anomalous patterns of mobility are also identified during evacuation and storm landfall, indicating how people respond to storm forecasts and early warnings.

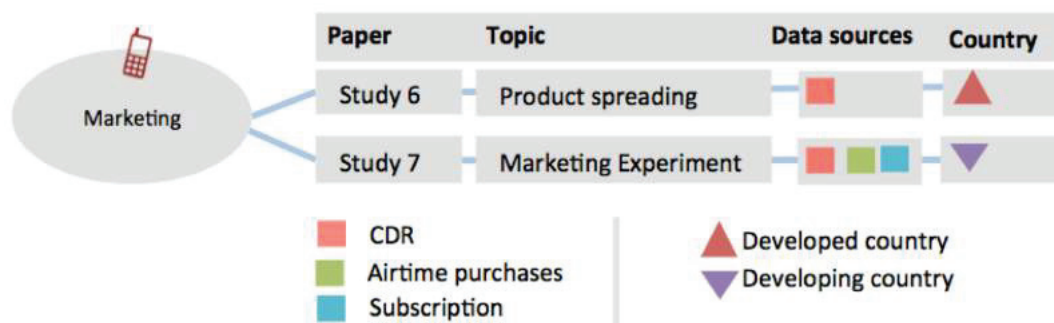
Detecting anomalous usage patterns from mobile network data is a promising avenue for researching human behavioural responses to environmental impacts across large spatiotemporal scales. Based on comparisons with rainfall measurements at landfall, and considering the considerable weakening of cyclones as they pass over land, calling frequency and population movement anomalies seem to be the best indicators of physical interaction and impact of the cyclone. The anomaly detection technique applied to CDRs, as presented here, overcomes some of these challenges, and demonstrates the potential value of CDR as a complement to current cyclone impact assessment tools to improve the accuracy, timeliness, and cost-effectiveness of cyclone impact assessments. Data from CDRs may be very useful as a tool to prioritise locations in which rapid needs assessment are performed after cyclone landfall, with the potential to drastically reduce the time to reach those most in need. Primary limitations of the study involve representativeness of the data for the general population. The indicators found should however reflect natural human response to shocks, but it is important

to stress that CDR data should not be used alone to indicate where post-disaster assistance is needed.

## 4.2 Contributions to behavioural insight and marketing

Building on the marketing literature from section 2.2.2 and the methodology introduced in section 3.1, the last two studies investigate how the social network matters when purchase decisions are made.

**Figure 4.8: Studies 6–7: topic, data sources and country**



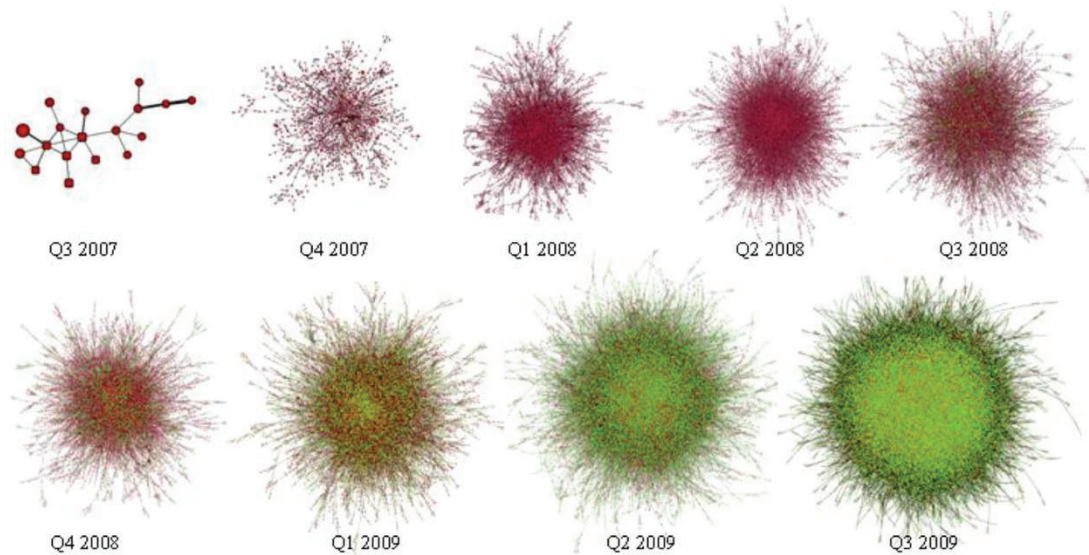
### Study 6: Product spreading over large-scale social networks

Most of the existing research on product diffusion in networks has been focused on a single product, with a static snapshot of the network. Here the interest is to know how the social network among product adopters develops over time. In addition the aim is to understand how the product diffusion depends on the underlying social network. In short this study looks for evidence of social spreading effects. The starting point of the analysis covers more than 2.5 million subscribers and 50 billion CDR records collected over 9 quarters.

By combining the social interactions, derived from CDRs, with the adoption history of a product of interest, how different products are spread over the social network can be observed. Several products have been studied, including the iPhone, iPad, video telephony, Doro and Android phones. For each product the evolution of fraction of adopters is shown by social component size, and centrality distributions by product is presented, as measured by eigenvector centrality of the adopters. All products show various degrees of social spreading effects. This work reveals in particular strong social effects in the adoption of iPhones and iPads. The Doro handsets have very weak social spreading, and will probably never ‘take

off". Video telephony also has strong spreading effects, but early takeoff was prevented by an external factor, here a new price model.


**Figure 4.9: Time evolution of the iPhone adoption network (largest-connected component). One node represents one subscriber. Node colour represents iPhone model: red = 2G, green = 3G, yellow = 3GS**



To verify social spreading effects the probability  $p_k$  that a subscriber has adopted a product is measured, given that  $k$  of the subscriber's friends have adopted the product. Figure 4.10 shows the following: if we know that a subscriber has one friend using a Doro phone, that subscriber's probability of using one themselves is roughly twice the adoption probability for a subscriber with no adopting friends.



**Figure 4.10: Adoption probability  $p_k$  vs. the number  $k$  of adopting friends for three products. In each case a monotonic growth of  $p_k$  with  $k$  is visible, indicating that some kind of social spreading is occurring**




Through visualisations, supported by analysis, a glimpse is provided of how products spread over large-scale social networks. These measurement methods give new and useful insight into how and why the different services have performed so differently. Initially social spreading effects were unexpected for Doro, which is a handset marketed towards the elderly. The findings quantify and suggest that such effects exist for most products. The strength of the effects also produce the possibility that such knowledge could be investigated further in marketing experiments.

### Study 7: Large-scale marketing using machine learning

Paper 6 quantified the adoption probability as a function of friends that have adopted a product. In paper 7 this insight is experimentally exploited in a campaign in Asia, using Internet packages as the product of choice. Economic growth in Asia is steadily increasing, but the overall Internet penetration is very small, which contributes to an economic and educational gap in comparison to more developed countries [93]. Mobile operators in this geographical region typically run thousands of text-based campaigns a year, resulting in customers receiving several promotional texts per month. This is a major problem for operators, since they are considered as spammers rather than providing people with useful information. For this particular operator, the policy is to not send more than one text per

customer every 14 days. It is therefore crucial to personalise the offers towards the customers. This paper evaluates a data-driven approach to marketing against the operator's current best practice in a large-scale 'Internet data' experiment. The experiment will compare the conversion rates of the treatment and control groups after one promotional text.

**Figure 4.11 a) Conversion rate in the control (best practice) and treatment (data-driven approach) groups; b) the percentage of converted people who renewed their data plan after using the volume included in the campaign offer**



The findings show that the data-driven approach using machine learning and SNA leads to higher conversion rates than the best-practice marketing approach.

The study also shows that historical natural adoption data can be used to train models when campaign response data is unavailable. New features to identify which customers are the most likely to convert into mobile Internet users were created. The features fall into the following categories: discretionary income, timing and social learning. In total 50 million customers were scored, and the treatment group was composed of the top 250,000 customers with the highest score. The conversion rate of an Internet data campaign was increased by 13 times compared to current best practice (the control group). The model also shows very good properties in the longer term, as 98% of the converted customers in the treatment group renewed their mobile internet packages after the campaign, compared to 37% in the control group.



**Table 4.1: Top 10 most predictive indicators to classify natural converters, ranked by importance in the model**

Rank	Type	Description
1	Social learning	Total spending on data among close social graph neighbours
2	Discretionary income	Average monthly spending on texts (binned)
3	Discretionary income	Average monthly number of texts sent (binned)
4	Discretionary income	Average monthly spending on value added services over text (binned)
5	Social learning	Average monthly spending on data among social graph neighbours
6		Data-enabled handset according to IMEI (Yes/No)
7	Social learning	Data volume among social graph neighbours
8	Social learning	Data volume among close social graph neighbours
9	Timing	Most used handset has changed since last month
10		Amount of 'accidental' data usage

The cross-validated model only relies on a few key variables. Only 20 features out of the initial 350 were selected for the final model. Table 4.1 shows the top 10 most useful features to classify natural converters.

The social learning features computed for this study turn out to be very helpful to help classify the natural converters. The total spending on data among the closest social graph neighbours is the most predictive feature. Using social features in selecting customers for the offer might have improved the retention rates. The results lead to speculation that the value customers derive from mobile data increases when their neighbours are also mobile data users. In other words, it is expected that a network externality effect exists in mobile Internet data. This means that selecting customers whose closest neighbours are already using mobile data might have locally used this network effect to create the lasting effect observed with a very high retention rate in the second month. It is also worth noting that the marketing team believed that average revenue per user (ARPU) would be the most predictive feature for Internet uptake. It can be seen however that the low-ARPU segment was slightly overrepresented, and the text- and data-focused discretionary spending variables are selected as being more important in the model. These findings open up exciting avenues of research within data-driven marketing and customer understanding. Such an approach has now been implemented in the respective business unit, and is greatly reducing spamming by providing the customer with more relevant offers. Such an approach can also be used to micro-target outreach or drive uptake of e.g. agricultural technologies, health-seeking behaviour or mobile money.

# 5 Discussion

Here the overall research question and challenges are addressed from a holistic perspective. The purpose is to analyse the contribution of this thesis, both in detail and from a societal perspective.

## 5.1 On the research question

**Q:** Apart from providing basic communication services, what kind of positive impact can we create for society and/or individuals using large-scale mobile phone datasets?

**A:** As shown in this thesis, new techniques could be adopted to achieve near-real-time insights into people's wellbeing and target aid interventions to vulnerable groups. The results show that mobile phone data can be used for making reliable socio-economic predictions on the individual and geographical level. The way you use a phone can also reveal something about 'who' you are, for example, whether you are illiterate or your income level. Such insight can further be aggregated to the geographical level to help vulnerable groups in society. Additional results show that mobile phone data can be combined with satellite information to estimate multi-dimensional poverty at high spatial resolution, using Bangladesh as an example. Eradication of poverty requires national and subnational quantification and monitoring, and existing poverty estimates rely on census data collected with limited temporal frequency. How traditional survey information may be complimented by adoption of this methodology has also been addressed.

Further results show that mobile phone data can be used to understand human behaviour during large systemic shocks and disasters in society. These results were exemplified by investigating the terror attack that hit Norway in 2011, and by analysing financial, social and mobility patterns from millions of people during a large cyclone event that hit Bangladesh in 2013. Such knowledge can be used to both gain understanding and detect early-warning signals that can help to prevent future disasters. Understanding how the diffusion of information occurs may also help us to understand the complex network in a modern society. In summary, these findings can be used to inform socially beneficial policies that include improved understanding of crisis behaviour, accurate mapping of service needs and faster

tracking and response. As obtaining survey information is costly and slow, there is a strong need for better approaches.

The following results are motivated by the questions of how people adopt new products and services, and what role the underlying social network structure plays in this process. Results indicate that social networks strongly matter when purchase decisions are made, and that these effects can be exploited to boost uptake of products and services. By including social patterns and machine learning techniques in a large-scale marketing experiment in Asia, the adoption rate was increased by 13 times compared to current best practice for a given product. This approach, which has now been adopted by the company, greatly reduces spam and provides customers with more relevant offers. At the same time the method provides an efficient way of increasing Internet penetration in a highly underdeveloped country. The same methodology may also be used for other services, as mentioned in section 4.2.2. In summary, these results provide additional insights into human behaviour with the aim to understand social interactions and improve marketing.

## **5.2 Challenges and limitations**

Using de-identified mobile phone logs represents many opportunities, but also throws up big challenges around public trust, privacy and potential abuses of data.

### **5.2.1 Privacy challenges**

Privacy challenges can be easy to overlook when confronted with the challenges of poverty, disease and basic economic growth. They are however critical to realise the great potential of innovations to help address these critical problems.

Privacy is defined by ITU as the ‘right of individuals to control or influence what information related to them may be disclosed’ [94]. Personal information that needs to be protected is central to the privacy framework. The OECD defines personal data as ‘any information relating to an identified or identifiable individual (data subject)’ [95]. The result of this approach has been the policy of ‘informed consent’ practiced by most companies to inform users of what data are being collected and how they are used. In the world of ‘Big Data’ this method has been argued to be impractical [96]. The main reason is that the greatest potential often lies in secondary applications long after the data has been collected. It is not possible

for companies to know a priori all the potential uses and to continuously seek permission from all users.

Other concerns relate to data anonymisation. Recent research on ‘computational privacy’ has shown that de-identification of users in CDR is possible, with high accuracy, if only four user data points are available [97]. In practice this is difficult as the real-world identities of users are unknown, but the authors argue that cross-referencing to other data sources might help to de-identify the individual. Such risks are currently lower in underdeveloped countries, as the amount of data stored on each user is still low. As addressed in papers 1 and 2, the great majority of mobile phone users in the developing world are prepaid, with minimal registration associated. In some countries, such as Pakistan, this might change, as mobile phone users must now go through a biometric verification before buying a SIM card. This creates a solid link between the owner of the SIM card and their actual identity [98].

One additional concern is when data go beyond metadata and move into the space of content information [99]. Linking non-personal data to an actual individual can then be made easier [100]. In particular this can be a problem when several sources of Big Data are coupled with other data sources to generate new insight.

### **5.2.2 Analysis and interpreting data**

Establishing a generalisation of findings from large-scale analysis might often be difficult, as several data sources are often involved and the chance of data quality issues increases. This section addresses some of the challenges related to analysis and data interpretation.

#### ***Data cleaning process***

Given multiple raw sources of data, it is important to understand how the data has mutated along the pathway into the final datasets. Some mobile phone operators choose to include the complete route of traffic that has been forwarded. This implies that there are multiple records in the CDRs for the same traffic sessions. The result is, for instance, errors when conducting SNA, where edge strengths are overvalued. Another concern is redundant locations, generated by the so-called ping-pong phenomenon where a user connects to several towers frequently. This might introduce errors when inferring the mobility patterns of users [101].

### *Selection bias and ground truth data*

A known concern is sampling selection bias. A large dataset may make the sampling rate irrelevant, but it does not make it representative [102]. The fact is that those who use a mobile phone are not necessarily a representative sample of the larger population being considered. This issue is particularly relevant when considering how mobile phone data may be used for monitoring, economic forecasting and development. Research studies are often based purely on data from one mobile operator, and depending on the type of data one can expect wealthier or poorer, and educated or uneducated individuals. Even if data from all operators in a country were available, nearing the total population, this is still not the whole population.

The main problem with results based on non-representative samples is that they lack external validity, which is the degree to which an internally valid conclusion can be generalised beyond the sample.

### *Change of behaviour and data sources*

As people adopt new technology, one challenge is the change of structure and content of mobile phone datasets. The signals of human behaviour in mobile phone datasets might change over time and place, analogously to stock markets where signals will change over time and new signals will appear or disappear. As communication signals might be transferred to other platforms, the CDR might not capture the social network as well as before [103]. The choice of data source(s) and when to re-train the models will be vital when considering the research questions. Another challenge is the change in human behaviour itself, which can be subject to self-censorship (turning the phone off in certain areas) and creation of multiple personas (SIM-switching).

### *Causation and correlation*

Big Data analysis draws much of its methodology from artificial intelligence and machine learning, which is primarily about prediction and correlations. Most often the datasets are observational and therefore are not able to measure causality. Hal Varian stated that ‘there are often more police in precincts with high crime, but that does not imply that increasing the number of police in a precinct would increase crime’ [14]. While some researchers have predicted the end of theory and hypotheses-testing [104], others are more sceptical [105]. Sendihl Mullainathan believes that inductive science (algorithmically mining data sources)

will not threaten deductive science (hypothesis testing). He believes that greater volume makes Big Data induction techniques more effective, while more variety makes them less effective. There will therefore be a need for explaining behaviour (deductive science) rather than just predicting behaviour. At the same time it is worth mentioning that causal modelling is possible in the Big Data regime by conducting large-scale experiments. Mobile phone operators utilise such techniques to learn about product usage and pricing.

### *The role of 'small' data*

Many of the publications in this thesis rely on a combination of survey data and large-scale datasets. To be able to create understanding and models that reflect ground realities, 'small' survey data will still be important in the future. For many studies it is crucial to verify the underlying assumptions in Big Data using survey data.

### **5.2.3 The future of artificial intelligence in social sciences**

Throughout this thesis machine learning and prediction methodology have been applied in a positive context, without much of a focus on potential negative pitfalls. Of course there are both short and long term implications of this emerging field that need to be considered from all angles. Artificial intelligence and machine learning are now being implemented in driverless cars [106], commercial drones [107], financial trading [108], autonomous weapons [109] and monitoring patients in care [110]. As such applications emerge it is important to be transparent around the decision-making process, especially since intelligent machines sometimes make errors too. In the field of social sciences this includes always validating the methodology to actual ground truth data, and using it as a complementary source of insight. It also includes updating analytical models, as researchers pointed out in relation to Google Flu trends, which made inaccurate forecasts for 100 out of 108 weeks [111]. Certainly, the development of artificial intelligence will increasingly make it a part of our daily lives. Currently we do not know if human-level intelligence will be incorporated into artificial intelligence, but it is important to consider it as a possible outcome of today's research.

## 6 Conclusion

This thesis has highlighted the economic and social benefits of using large-scale mobile phone metadata. This data paints a picture of an individual's patterns of behaviour and interactions. When applying this capability across a nation covering millions of people, it is possible to collapse a very detailed picture of the entire population. There are still many challenges to overcome, especially when addressing the privacy implications of Big Data, but mobile phone records can be mutually beneficial to both the private sector and the government.

This thesis has addressed how this information can be used to inform socially beneficial policies on socioeconomics, poverty and disasters on the society level. More targeted and timely policy actions can be taken by governments to address the underlying problems, which would not be possible with the limited and lagged insights revealed by traditional official statistics.

Furthermore, the possibility of using new behavioural insights has been explored, by exploiting large-scale social networks to build models and target the right individuals in marketing experiments. By incorporating social information, product adoption is significantly increased and spam is reduced.

Hopefully, this dissertation will contribute to generating debates and interest among a wide range of audiences, to advancing the understanding of analysis of large-scale datasets and their applications, and most importantly to making such analyses possible and beneficial for society.

# Bibliography

- [1] L. Einav and J. Levin, "Economics in the age of big data," *Science*, vol. 346, no. 6210, p. 1243089, 2014.
- [2] OECD, "Data-Driven Innovation: Big Data for Growth and Well-Being," Paris, Technical report 2015.
- [3] A. De Mauro, M. Greco, and M. Grimaldi, "What is big data? A consensual definition and a review of key research topics," in *AIP Conference Proceedings*, vol. 1644, 2015, pp. 97-104.
- [4] D.A. Kim et al., "Social network targeting to maximise population behaviour change: A cluster randomised controlled trial," *The Lancet*, vol. 386, no. 9989, pp. 145-153, 2015.
- [5] S. Aral and D. Walker, "Identifying influential and susceptible members of social networks," *Science*, vol. 337, no. 6092, pp. 337-341, 2012.
- [6] J. Blumenstock, G. Cadamuro, and R. On, "Predicting poverty and wealth from mobile phone metadata," *Science*, vol. 350, no. 6264, pp. 1073-1076, 2015.
- [7] N. Jean et al., "Combining satellite imagery and machine learning to predict poverty," *Science*, vol. 353, no. 6301, pp. 790-794, 2016.
- [8] M.C. Gonzalez, C.A. Hidalgo, and A.L. Barabasi, "Understanding individual human mobility patterns," *Nature*, vol. 453, no. 7196, pp. 779-782, 2008.
- [9] S. Aral and D. Walker, "Creating social contagion through viral product design: A randomized trial of peer influence in networks," *Management Science*, vol. 57, no. 9, pp. 1623-1639, 2011.
- [10] A. Wesolowski et al., "Quantifying the impact of human mobility on malaria," *Science*, vol. 338, no. 6104, pp. 267-270, 2012.
- [11] UN Global Pulse, "Big data for development: Challenges & opportunities," New York, White paper 2012.
- [12] D. Lazer et al., "Life in the network: the coming age of computational social science," *Science*, vol. 323, no. 5915, p. 721, 2009.
- [13] E. Adar, "The two cultures and Big Data Research," *Journal of Law and Policy for the Information Society*, vol. 10, no. 3, pp. 765-781, 2015.
- [14] H. Varian, "Big Data: New Tricks for Econometrics," *Journal of Economic Perspectives*, vol. 28, no. 2, pp. 3-27, 2014.
- [15] J. Kleinberg, J. Ludwig, S. Mullainathan, and Z. Obermeyer, "Prediction Policy Problems," *The American economic review*, vol. 105, no. 5, pp. 491-495, 2015.
- [16] GSMA, "GSMA Global Economy Report," Technical report 2015.
- [17] R. Horak, *Telecommunications and data communications handbook*.: John Wiley & Sons, 2007.
- [18] N. Hughes and S. Lonie, "M-PESA: mobile money for the "unbanked" turning cellphones into 24-hour tellers in Kenya," *Innovations*, vol. 2, no. 1-2, pp. 63-81, 2013.
- [19] V.D. Blondel et al., "Data for development: the d4d challenge on mobile phone data," *arXiv preprint:1210.0137*, 2012.
- [20] Y.A. de Montjoye, Z. Smoreda, R. Trinquart, C. Ziemlicki, and V.D. Blondel, "D4D-Senegal: the second mobile phone data for development challenge," *arXiv preprint:1407.4885*, 2014.
- [21] K. Kumar and K. Muhota, "Can Digital Footprints Lead to Greater Financial Inclusion?," The World Bank, No 9439 2012.



- [22] N. Eagle, A.S. Pentland, and D. Lazer, "Inferring friendship network structure by using mobile phone data," *Proceedings of the national academy of sciences*, vol. 106, no. 36, pp. 15274-15278, 2009.
- [23] C. Riedl et al., "Diffusion of Information Through On-Demand Information Seeking Behavior," *Forthcoming paper*, 2017.
- [24] B.C. Csáji et al., "Exploring the mobility of mobile phone users," *Physica A: statistical mechanics and its applications*, vol. 392, no. 6, pp. 1459-1473, 2015.
- [25] S. Phithakkitnukoon, Z. Smoreda, and P. Olivier, "Socio-geography of human mobility: A study using longitudinal mobile phone data," *PloS ONE*, vol. 7, no. 6, p. e39253, 2012.
- [26] M. Tizzoni et al., "On the use of human mobility proxies for modeling epidemics," *PLoS Comput Biol*, vol. 10, no. 7, p. e1003716, 2014.
- [27] J. Rana et al., "Handset-centric view of smartphone application use," *Procedia Computer Science*, vol. 34, pp. 368-375, 2014.
- [28] White House report, "Big Data: Seizing opportunities, preserving value," Washington, DC: Executive Office of the President., 2014.
- [29] N. Eagle, M. Macy, and R. Claxton, "Network diversity and economic development," *Science*, vol. 328, no. 5981, pp. 1029-1031, 2010.
- [30] E. Jahani et al., "Networks and Income: Evidence from Individually Matched Income and Mobile Phone Metadata," in *Forthcoming paper*, 2017.
- [31] C. Smith, A. Mashhadi, and L. Capra, "Ubiquitous sensing for mapping poverty in developing countries," in *3rd International Conference on the Analysis of Mobile Phone Datasets*, 2013.
- [32] V. Frias-Martinez and J. Virseda, "On the relationship between socio-economic factors and cell phone usage," in *Proceedings of the fifth international conference on information and communication technologies and development*, 2012, pp. 76-84.
- [33] T. Gutierrez, G. Krings, and V.D. Blondel, "Evaluating socio-economic state of a country analyzing airtime credit and mobile phone datasets," in *arXiv preprint:1309.4496*, 2013.
- [34] A. Decuyper et al., "Estimating Food Consumption and Poverty indices with Mobile Phone Data," in *arXiv preprint:1412.2595*, 2014.
- [35] P. Deville et al., "Dynamic population mapping using mobile phone data," *Proceedings of the National Academy of Sciences*, vol. 111, no. 45, pp. 15888-15893, 2014.
- [36] H. Sterly, B. Hennig, and K. Dongo, "'Calling abidjan' - improving population estimations with mobile communication data," in *Improving Population Estimations with Mobile Communication Data*, 2013.
- [37] K. Tinsley and M. Bishop, "Poverty and Population Density: Implications for Economic Development Policy," *Journal of Higher Education Outreach and Engagement*, vol. 11, no. 1, pp. 195-208, 2007.
- [38] J.L. Toole et al., "Tracking employment shocks using mobile phone data," *The Royal Society Interface*, vol. 12, no. 107, p. 20150185, 2015.
- [39] D. Brookshire and M. McKee, "Other indirect costs and losses from earthquakes: issues and estimation," *Indirect Economic Consequences of a Catastrophic Earthquake*, pp. 267-325, 1992.
- [40] J.P. Bagrow, D. Wang, and A.L. Barabasi, "Collective response of human populations to large-scale emergencies," *PloS one*, vol. 6, no. 3, p. 17680, 2011.
- [41] L. Bengtsson, X. Lu, A. Thorson, R. Garfield, and J. Von Schreeb, "Improved response to disasters and outbreaks by tracking population movements with mobile phone network data: a post-earthquake geospatial study in Haiti," *PLoS Med*, vol. 8, no. 8, p. e1001083, 2011.
- [42] X. Lu, L. Bengtsson, and P. Holme, "Predictability of population displacement after the

- 2010 Haiti earthquake," *Proceedings of the National Academy of Sciences*, vol. 109, no. 29, pp. 11576-11581, 2012.
- [43] T. Philipson, "Economic Epidemiology and Infectious Diseases," *Handbook of health economics*, vol. 1, pp. 1761-1799, 2000.
  - [44] A. Wesolowski et al., "Commentary: Containing the ebola outbreak—the potential and challenge of mobile network data," *PLOS Currents Outbreaks*, 2014.
  - [45] A. Wesolowski et al., "Impact of human mobility on the emergence of dengue epidemics in Pakistan," *Proceedings of the National Academy of Sciences*, vol. 112, no. 38, pp. 11887-11892, 2015.
  - [46] A. Banerjee, E. Duflo, and N. Qian, "On the Road: Access to Transportation Infrastructure and Economic Growth in China," *NBER Working Paper No. 17897*, 2012.
  - [47] M. Berlingerio, F. Calabrese, G. DiLorenzo, N. R. Nair, and F. Pinelli, "AllAboard: A system for exploring urban mobility and optimizing public transport using cellphone data," in *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, 2013, pp. 663-666.
  - [48] J.L. Toole, S. Colak, F. Alhasoun, A. Evsukoff, and M.C. Gonzalez, "The path most travelled: Mining road usage patterns from massive call data.," *arXiv preprint:1403.0636*, 2014.
  - [49] S. Çolak, A. Lima, and M.C. González, "Understanding congested travel in urban areas," *Nature Communications*, vol. 7, 2016.
  - [50] J. Steenbruggen, M.T. Borzacchiello, P. Nijkamp, and H. Scholten, "Mobile phone data from gsm networks for traffic parameter and urban spatial pattern assessment: a review of applications and opportunities.," *GeoJournal*, vol. 78, no. 2, pp. 223-243, 2013.
  - [51] S. Aral, "Identifying Social Influence: A Comment on Opinion Leadership and Social Contagion in New Product Diffusion," *Marketing Science*, vol. 30, no. 2, pp. 217-223, 2011.
  - [52] A. Sundararajan, F. Provost, G. Oestreicher-Singer, and S. Aral, "Research commentary-information in digital, economic, and social networks," *Information Systems Research*, vol. 24, no. 4, pp. 883-905, 2013.
  - [53] M. Trusov, A.V. Bodapati, and R.E. Bucklin, "Determining Influential Users in Internet Social Networks," *Journal of Marketing Research*, vol. 47, no. 4, pp. 643-658, 2010.
  - [54] E. Katz and P.F. Lazarsfeld, "Personal Influence, The part played by people in the flow of mass communications," *Transaction Publishers*, 1966.
  - [55] S. Hill, F. Provost, and C. Volinsky, "Network-based marketing: Identifying likely adopters via consumer networks," *Statistical Science*, pp. 256-276, 2006.
  - [56] N. Aharony, W. Pan, C. Ip, and A. Pentland, "Tracing mobile phone app installations in the "friends and family" study," in *Proceedings of the Workshop on Information in Networks (WIN)*, 2010.
  - [57] H. Risselada, P.C. Verhoef, and T.H. Bijmolt, "Dynamic effects of social influence and direct marketing on the adoption of high-technology products.," *Journal of Marketing*, vol. 78, no. 2, pp. 52-68, 2014.
  - [58] I. Mas and O. Morawczynski, "Designing mobile money services lessons from M-PESA," *Innovations*, vol. 4, no. 2, pp. 77-91, 2009.
  - [59] M.R. Khan and J.E. Blumenstock, "Predictors without Borders: Behavioral Modeling of Product Adoption in Three Developing Countries," in *Knowledge Discovery & Data Mining (KDD)*, 2016.
  - [60] Telenor Group. (2016) Telenor Group Operations. [Online]. <https://www.telenor.com/>
  - [61] C.J. Date and H. Darwen, *A Guide To Sql Standard.*: Addison-Wesley, 1997.
  - [62] G. Van Rossum, "Python Programming Language," in *USENIX Annual Technical Conference*, vol. 41, 2007.
  - [63] R. Ihaka and R. Gentleman, "R: a language for data analysis and graphics," *Journal of*

- Computational and Graphical Statistics*, vol. 5, no. 3, pp. 299-314, 1993.
- [64] Autodesk. (1998) Autodesk Maya. [Online]. <http://www.autodesk.com/>
  - [65] P. Shannon, A. Markiel, O. Ozier, and N.S. Baliga, "Cytoscape: a software environment for integrated models of biomolecular interaction networks," *Genome research*, vol. 13, no. 11, pp. 2498-2504, 2003.
  - [66] QGIS. (2002) QGIS. [Online]. <http://www.qgis.org/>
  - [67] H. Rue, S. Martino, and N. Chopin, "Approximate Bayesian inference for latent Gaussian models by using integrated nested Laplace approximations," *Journal of the royal statistical society: Series b (statistical methodology)*, vol. 71, no. 2, pp. 319-392, 2009.
  - [68] S.P. Borgatti, A. Mehra, D.J. Brass, and G. Labianca, "Network Analysis in the Social Sciences," *Science*, vol. 323, no. 5916, pp. 892-895, 2009.
  - [69] M. Newman, *Networks: an introduction.*: Oxford University Press, 2010.
  - [70] M. Newman, "The mathematics of networks," *The new palgrave encyclopedia of economics*, vol. 2, pp. 1-12, 2008.
  - [71] J.P. Onnela et al., "Structure and tie strengths in mobile communication networks," *Proceedings of the National Academy of Sciences*, vol. 104, no. 18, pp. 7332-7336, 2007.
  - [72] C.M. Bishop, "Pattern recognition," *Machine Learning*, vol. 128, 2006.
  - [73] G.H. Givens and J.A. Hoeting, *Computational statistics.*: John Wiley & Sons, 2012, vol. 710.
  - [74] C. Shearer, "The CRISP-DM model: the new blueprint for data mining," *Journal of data warehousing*, vol. 5, no. 4, pp. 13-22, 2000.
  - [75] R. Kohavi, "A study of cross-validation and bootstrap for accuracy estimation and model selection," *International Joint Conference on Artificial Intelligence*, vol. 14, no. 2, pp. 1137-1145, 1995.
  - [76] J.A. Hanley and B.J. McNeil, "The meaning and use of the area under a receiver operating characteristic (ROC) curve," *Radiology*, vol. 143, no. 1, pp. 29-36, 1982.
  - [77] T. Hastie, R. Tibshirani, and J. Friedman, "Supervised learning," in *The Elements of Statistical Learning*. New York: Springer, 2005, pp. 9-484.
  - [78] E Alpaydin, *Introduction to machine learning.*: MIT press, 2014.
  - [79] Y. LeCun, Y. Bengio, and G. Hinton, "Deep Learning," *Nature*, vol. 521, no. 7553, pp. 436-444, 2015.
  - [80] M. Mohri, A. Rostamizadeh, and A. Talwalkar, *Foundations of machine learning.*: MIT press, 2012.
  - [81] R. Caruana and A. Niculescu-Mizil, "An empirical comparison of supervised learning algorithms," in *Proceedings of the 23rd international conference on Machine learning*, 2006, pp. 161-168.
  - [82] T. Hastie, R. Tibshirani, and J. Friedman, "Unsupervised learning," in *The Elements of Statistical Learning.*: Springer New York, 2005, pp. 485-585.
  - [83] J. Grimmer, "We are all social scientists now: how big data, machine learning, and causal inference work together," *Political Science & Politics*, vol. 48, no. 1, pp. 80-83, 2015.
  - [84] S. Athey and G.W. Imbens, "Machine learning methods for estimating heterogeneous causal effects," *stat*, vol. 5, 2015.
  - [85] A. Cree, A. Kay, and J. Steward, "The economic and social cost of illiteracy: A snapshot of illiteracy in a global context," World Literacy Foundation, Report 2012.
  - [86] A. Krizhevsky, I. Sutskever, and G.E. Hinton, "Imagenet classification with deep convolutional neural networks," *Advances in neural information processing systems*, pp. 1097-1105, 2012.
  - [87] G.E. Dahl, D. Yu, L. Deng, and A. Acero, "Context-dependent pre-trained deep neural networks for large-vocabulary speech recognition," *IEEE Transactions on Audio, Speech*,

- and *Language Processing*, vol. 20, no. 1, pp. 30-42, 2012.
- [88] M. Cruz, J. Foster, B. Quillin, and P. Schellekens, "Ending extreme poverty and sharing prosperity: Progress and policies," World Bank Group, Policy Research Note 2015.
  - [89] D. Griggs et al., "Policy: Sustainable development goals for people and planet," *Nature*, vol. 495, no. 7441, pp. 305-307, 2013.
  - [90] W. Dutton and F. Nainoa, "The Social Dynamics of Wireless on September 11 : Reconfiguring Access," *Crisis Communication*, pp. 69-82, 2003.
  - [91] A.A. Cohen and D. Lemish, "When bombs go off the mobiles ring: The aftermath of terrorist attacks," in *A sense of place: The global and the local in mobile communication*. Vienna: Passagen Verlag, 2005, pp. 117-128.
  - [92] M. Burke, S.M. Hsiang, and E. Miguel, "Global non-linear effect of temperature on economic production," *Nature*, vol. 527, pp. 235-239, 2015.
  - [93] M.D. Chinn and R.W. Fairlie, "The determinants of the global digital divide: a cross-country analysis of computer and internet penetration," *Oxford Economic Papers*, vol. 59, no. 1, pp. 16-44, 2007.
  - [94] ITU, "Security in Telecommunications and Information Technology: An overview of issues and the deployment of existing ITU-T Recommendations for secure telecommunications," Geneva, Report 2003.
  - [95] OECD, "The OECD Privacy Framework," Technical report 2013.
  - [96] C. Kalapesi, "Unlocking the value of personal data: From collection to usage," World Economic Forum, Technical report 2013.
  - [97] Y.A. de Montjoye, C.A. Hidalgo, M. Verleysen, and V.D. Blondel, "Unique in the Crowd: The privacy bounds of human mobility," *Scientific reports*, vol. 3, 2013.
  - [98] J. Bigün, J. Fierrez-Aguilar, J. Ortega-Garcia, and J. Gonzalez-Rodriguez, "Multimodal biometric authentication using quality signals in mobile communications," in *International Conference on Image Analysis and Processing*, 2003, pp. 2-11.
  - [99] M. Kosinski, D. Stillwell, and T. Graepel, "Private traits and attributes are predictable from digital records of human behavior," *Proceedings of the National Academy of Sciences*, vol. 110, no. 15, pp. 5802-5805, 2013.
  - [100] P. Ohm, "Broken promises of privacy: Responding to the surprising failure of anonymization," *UCLA law review*, vol. 57, p. 1701, 2010.
  - [101] C. Iovan, A.M. Olteanu-Raimond, T. Couronné, and Z. Smoreda, "Moving and Calling: Mobile Phone Data Quality Measurements and Spatiotemporal Uncertainty in Human Mobility Studies," in *Geographic Information Science at the Heart of Europe.*: Springer International Publishing, 2013, pp. 247-265.
  - [102] D. Boyd and K. Crawford, "Critical questions for big data: Provocations for a cultural, technological, and scholarly phenomenon," *Information, communication & society*, vol. 15, no. 5, pp. 662-679, 2012.
  - [103] K. Church and R. de Oliveira, "What's up with whatsapp?: comparing mobile instant messaging behaviors with traditional SMS," in *Proceedings of the 15th international conference on Human-computer interaction with mobile devices and services*, 2013, pp. 352-361.
  - [104] C. Anderson, "The end of theory: the data deluge makes the scientific method," *Wired*, 2008.
  - [105] S. Mullainathan, "What Big Data Means For Social Science," HeadCon'13, 2013.
  - [106] T. Luettel, M. Himmelsbach, and H.J. Wuensche, "Autonomous ground vehicles—Concepts and a path to the future," in *Proceedings of the IEEE*, vol. 100, 2012, pp. 1831-1839.
  - [107] T. Krajník, V. Fišer, D. Vonásek, and J. Faigl, "AR-drone as a platform for robotic research and education," *Robotics*, pp. 172-186, 2011.

- [108] J.B Heaton, J.H. Polson, and J.H. Witte, "Deep Learning in Finance," *arXiv preprint:1602.06561*, 2016.
- [109] P. Asaro, "On banning autonomous weapon systems: human rights, automation, and the dehumanization of lethal decision-making," *International Review of the Red Cross*, vol. 94, no. 886, pp. 687-709, 2012.
- [110] C.R. Farrar and K. Worden, *Structural health monitoring: a machine learning perspective.*: John Wiley & Sons, 2012.
- [111] D. Lazer, R. Kennedy, G. King, and A. Vespignani, "The parable of Google flu: traps in big data analysis," *Science*, vol. 343, no. 6176, pp. 1203-1205, 2014.



# Publications

## **Paper 1**

Can mobile usage predict illiteracy in a developing country?..

## **Paper 2**

Deep learning applied to mobile phone data for Individual income classification

## **Paper 3**

Mapping Poverty using mobile phone and satellite data

## **Paper 4**

The activation of core social networks in the wake of the 22 July Oslo bombing

## **Paper 5**

Detecting climate adaptation with mobile network data: Anomalies in communication, mobility and consumption patterns during Cyclone Mahasen

## **Paper 6**

Comparing and visualizing the social spreading of products on a large-scale social network

## **Paper 7**

Big Data-Driven Marketing: How Machine Learning outperforms marketers' gut-feeling





# **Paper 1**

---



## **Paper 2**

---



# Deep learning applied to mobile phone data for Individual income classification

Pål Sundsøy<sup>1\*</sup>, Johannes Bjelland<sup>1</sup>, Bjørn-Atle Reme<sup>1</sup>, Asif M.Iqbal<sup>1</sup> and Eaman Jahani<sup>2</sup>

<sup>1</sup>Telenor Group Research, Snarøyveien 30,1331 Fornebu, Norway

<sup>2</sup>Massachusetts Institute of Technology, The Media Laboratory, USA

\*pal-roe.sundsoy@telenor.com

**Abstract—** Deep learning has in recent years brought breakthroughs in several domains, most notably voice and image recognition. In this work we extend deep learning into a new application domain - namely classification on mobile phone datasets. Classic machine learning methods have produced good results in telecom prediction tasks, but are underutilized due to resource-intensive and domain-specific feature engineering. Moreover, traditional machine learning algorithms require separate feature engineering in different countries. In this work, we show how socio-economic status in large de-identified mobile phone datasets can be accurately classified using deep learning, thus avoiding the cumbersome and manual feature engineering process. We implement a simple deep learning architecture and compare it with traditional data mining models as our benchmarks. On average our model achieves 77% AUC on test data using location traces as the sole input. In contrast, the benchmarked state-of-the-art data mining models include various feature categories such as basic phone usage, top-up pattern, handset type, social network structure and individual mobility. The traditional machine learning models achieve 72% AUC in the best-case scenario. We believe these results are encouraging since average regional household income is an important input to a wide range of economic policies. In underdeveloped countries reliable statistics of income is often lacking, not frequently updated, and is rarely fine-grained to sub-regions of the country. Making income prediction simpler and more efficient can be of great help to policy makers and charity organizations – which will ultimately benefit the poor.

*Deep learning; Mobile phone data; poverty;household income ; Big Data Analytics; Machine learning; Asia, Mobile network operator; metadata; algorithms*

## I. INTRODUCTION

Recent advances in Deep Learning [1][2] have made it possible to extract high-level features from raw sensory data, leading to breakthroughs in computer vision [9][10][11] and speech recognition [12][13]. It seems natural to ask whether similar techniques could also be beneficial for useful prediction tasks on mobile phone data, where classic machine learning algorithms are often under-utilized due to time-consuming country and domain- specific feature engineering [6].

Our work investigates how we can separate individuals with high and low socio-economic status using mobile phone call detail records (CDR). Finding good proxies for income in mobile phone data could lead to better poverty prediction – which could ultimately lead to more efficient policies for addressing extreme poverty in hardest hit regions.

With this in mind, we perform a large-scale country-representative survey in a low HDI Asian country, where household income is collected from approximately 80 000 individuals. The individual records are de-identified and coupled with raw mobile phone data that span over 3 months. This dataset allow us to build a deep learning model, as well as a benchmarking model using custom feature engineering and traditional data mining algorithms.

From the household income we derive two binary classifiers (1) below or above median household income and (2) below or above upper poverty level. The income threshold for poverty level is derived from official statistics and based on average national household income and household size. Our survey classifies participants into 13 household income bins – where bin 1 and 2 correspond to below upper poverty level.

The rest of this paper is organized as follows: In section 2 we describe the features and models used for benchmarking our deep learning approach. Section 3 describes the deep learning approach itself. In section 4 we compare the results of the two approaches. Finally, we draw our conclusions in section 5.

## II. BEST PRACTICE – TRADITIONAL DATA MINING

This section describes the features and the standard machine learning algorithms used as our benchmark. The data preparation phase in the data mining process is a tedious task, that requires specific knowledge about both the local market and the various data channels as potential sources for input features. Typically the data warehouse architecture and the data format vary between the operators and third-party software packages, making it hard to create generalizable feature-sets.

### A. Features

We build a structured dataset consisting of 150 features from 7 different feature families, see Table 1. The features are

custom-made and range from basic phone usage, top-up pattern, social network and mobility to handset usage. The features include various parameters of the corresponding distributions such as weekly or monthly median, mean and variance.

**TABLE 1** SAMPLE FEATURES USED IN TRADITIONAL MACHINE LEARNING BENCHMARKS.

Feature family	Feature examples
Basic phone usage	Outgoing/incoming voice duration, sms count etc.
Top-up transactions	Spending speed, recharge amount per transaction, fraction of lowest/highest recharge amount, coefficient of variation recharge amount etc
Location/mobility	Home district/tower, radius of gyration, entropy of places, number of places etc.
Social Network	Interaction per contact, degree, entropy of contacts etc.
Handset type	Brand, manufacturer, camera enabled, smart/feature/basic phone etc
Revenue	Charge of outgoing/incoming SMS, MMS, voice, video, value added services, roaming, internet etc.
Advanced phone usage	Internet volume/count, MMS count, video count/duration, value added services duration/count etc.

#### A. Models

Ensemble methods have been proven as powerful algorithms when applied to large-scale telecom datasets [14][6]. They combine predictions of several base classifiers built with a given learning algorithm in order to improve robustness over a single classifier. We investigate two different state-of-the-art ensemble methods:

1) *Random forest*, where we build several independent classifiers and average their predictions, thus reducing the variance. We choose grid search to optimize the tree size.

2) *Gradient boosting machines* (GBM) where the base classifiers are built sequentially. The algorithm combines the new classifier with ones from previous iterations in an attempt to reduce the overall error rate. The main motivation is to combine several weak models to produce a powerful ensemble.

In our set-up, each model is trained and tested using a 75/25 split. For the response variable related to poverty level we introduce misclassification cost. Since few people are below the poverty level (minority class), a naive model will predict everyone above poverty line. We therefore apply a misclassification cost to the minority class to achieve fewer false positives and adjust for the ratio between the classes.

### III. DEEP LEARNING

In this section we describe the input data and the structure of our deep learning model.


#### B. Features

Classification results of the traditional learning algorithms are inherently limited in performance by the quality of the extracted features [7]. Deep learning can instead reproduce complicated functions that represent higher level extractions, and replace manual domain-specific feature engineering.

Earlier studies have shown a good correlation between location/human mobility and socio-economic levels [15-17]. Using this as a motivation we build a simple vector whose length corresponds to the number of mobile towers (8100 dimensions), and the vector elements correspond to the mobile phone activity at the given tower – shown in Table 2.

**TABLE 2** INPUT VECTORS TO DL ALGORITHM BEFORE NORMALIZATION

Hashed Phone number	Tower 1	Tower2	....	Tower8100	Below/above poverty level
1	0	67	....	16	0
2	7	0	....	0	1
.	.	.	....	.	.
.	.	.	....	.	.
80K	0	9	....	6	0




**FIGURE 1** LOCATION ACTIVITY MAP. THE MAP PROVIDES INFORMATION ABOUT EACH USERS' TOWER ACTIVITY.

A visual representation of Table 2 can be seen in Fig 1, where each subscriber is represented by a location density map.

#### A. Models

We use a standard multi-layer feedforward architecture where the weighted combination of the  $n$  input signals is aggregated, and an output signal  $f(\alpha)$  is transmitted by the connected neuron. The function  $f$  used for the nonlinear

activation is rectifier  $f(\alpha) \approx \log(1+e^\alpha)$ . To minimize the loss function we apply a standard stochastic gradient descent with the gradient computed via back-propagation. We use dropout [4][5] as a regularization technique to prevent over-fitting. For the input layer we use the value of 0.1 and 0.2 for the hidden layers. Dropout secures that each training example is used in a different model which all share the same global parameters. This allows the models to be averaged as ensembles, improving generalization and preventing overfitting. The split between train and test set, as well as the introduced misclassification cost for poverty level, are similar to the benchmark models.



**FIGURE 2** AUC ON TEST SET SHOWING THE TRUE POSITIVE VS FALSE POSITIVE RATE FOR DEEP LEARNING (DL), GRADIENT BOOSTING MACHINES (GBM) AND RANDOM FOREST (RF).


**TABLE 3** MODEL PERFORMANCE (AUC) ON TEST SET

Model	AUC test set	
	Below/above poverty level	Below/above median income
DL	<b>0.74</b>	<b>0.77</b>
GBM	0.68	0.71
RF	0.69	0.72

#### IV. RESULTS

Our deep neural network is trained to separate individuals with high and low socio-economic status. We evaluate our models' performance with AUC [2][3] achieving 77% and 74% AUC on test set when the classifier is predicting above/below median household income and above/below poverty level respectively (Fig 2, Table 3). The corresponding AUC on the train sets are respectively 80% and 77%, showing that our model does not overfit significantly. Our results indicate that the DL model achieves a higher performance than those of multi-source RF and GBM models, which vary between 71-72% and 68-69% for above/below median household income and poverty level respectively.

Next, we feed the RF and GBM models with the same representation as fed into the DL algorithm - achieving RF performance of AUC 64% and GBM performance of AUC 61%. The performance of these models greatly suffers as the number of input features increase without increasing the training size. We conclude that given a fixed training sample, the traditional models are not able to learn a complicated function that represents higher level extractions, but perform better when using manual domain-specific feature engineering.



**FIGURE 3** TOP FEATURES IN RF MODEL COLORED BY THEIR RESPECTIVE FEATURE FAMILY.

A posteriori inspection of the top features selected by our traditional models leads to some interesting qualitative insights. We observe a similar pattern in the top features selected by the RF and GBM model. Figure 3 shows the top features in our RF model. We notice the importance of three feature families:

1) *Location dynamics*: Where the user spends most of his time is a good signal of his income. This indicates that our models have detected regions of low economic development status.

2) *The handset brand*: In the country of our study, minimal and more affordable handset brands are very popular among the lower income quantiles, while expensive smartphones are considered as a huge status symbol.

3) *The top-up pattern*: Interestingly, the recharge amount per transaction is more predictive than the total recharge amount. We observe that individuals from the lower income quantiles usually top-up with lower amounts when they first fill up their account.

## V. CONCLUSION

In order to predict household income based on mobile phone communication and mobility patterns, we implemented a multi-layer feedforward deep learning architecture. Our approach introduces a novel data representation for learning neural networks on real CDR data.

Our approach suggests that multi-layer feedforward models are an effective tool for predicting economic indicators based on mobile communication patterns. While capturing the complex dependencies between different dimensions of the data, deep learning algorithms do not overfit the training data as seen by our test performance.

Furthermore, our deep learning model, using only a single dimension of the data in its raw form, achieves a 7% better performance compared to the best traditional data mining approach based on custom engineered features from multiple data dimensions. Even though such an automated approach is time-saving, many of the classic machine learning approaches have the advantage of being interpretable. However, since a large portion of a data mining process is data preparation, there is a big demand to automate this initial step.

As future work, we would like to investigate the performance implications of including temporal aspects of raw CDRs in our models and the data representation. In addition, we will work on finding a general representation of telecom data that can be used for various prediction tasks.

An important application of this work is the prediction of regional and individual poverty levels in low HDI countries. Since our approach only requires de-identified customer and tower IDs, we find this method more privacy preserving compared to traditional data mining approaches where the input features may reveal sensitive information about the customers.

## REFERENCES

- [1] Yann LeCun, Yoshua Bengio & Geoffrey Hinton. Deep Learning, Nature 521 436-444 (28 May 2015)
- [2] P.Baldi et al. Searching for exotic particles in high-energy physics with deep learning, Nature Communications 5, article 4308
- [3] AUC: A better measure than accuracy in comparing learning algorithms, Ling C, Advances in Artificial Intelligence LNCS Vol 2671, 2003, pp 329-341
- [4] G. Hinton et. al. (2012) Improving Neural Networks by Preventing Co-adaptation of Feature Detectors. arXiv preprint arXiv:1207.0580
- [5] G. Dahl et. al. (2013) Improving Deep Neural Networks for LVCSR using Rectified Linear Units and Dropout. ICASSP, 8609-8613.
- [6] P.Sundsoy, J.Bjelland, A.Iqbal, A.Pentland, Y.A.Montjoye, Big Data-Driven Marketing: How machine learning outperforms marketers' gut-feeling, Social Computing, Behavioral-Cultural Modeling & Prediction. Lecture Notes in Computer Science Volume 8393, 2014, pp 367-374
- [7] Arel I, Rose DC, Karnowski TP, Deep machine learning – A new frontier in artificial intelligence research. IEEE computational intelligence magazine 5, 13-18 (2010).
- [8] Bengio Y, Learning deep architectures for AI. Foundations and trends in machine learning 2 (2009).
- [9] Alex Krizhevsky, Ilya Sutskever, and Geoff Hinton. Imagenet classification with deep convolutional neural networks. In Advances in Neural Information Processing Systems 25, pages 1106–1114, 2012
- [10] Pierre Sermanet, Koray Kavukcuoglu, Soumith Chintala, and Yann LeCun. Pedestrian detection with unsupervised multi-stage feature learning. In Proc. International Conference on Computer Vision and Pattern Recognition (CVPR 2013). IEEE, 2013.
- [11] Volodymyr Mnih. Machine Learning for Aerial Image Labeling. PhD thesis, University of Toronto, 2013.
- [12] George E. Dahl, Dong Yu, Li Deng, and Alex Acero. Context-dependent pre-trained deep neural networks for large-vocabulary speech recognition. Audio, Speech, and Language Processing, IEEE Transactions on, 20(1):30–42, January 2012.
- [13] Alex Graves, Abdel-rahman Mohamed, and Geoffrey E. Hinton. Speech recognition with deep recurrent neural networks. In Proc. ICASSP, 2013
- [14] Namyoung K, Kyu J, Yong K, A New Ensemble Model for Efficient Churn Prediction in Mobile Telecommunication, 2012 45<sup>th</sup> Hawaii international conf on system sciences
- [15] V.Frias-Martinez, J.Virseda, E.Frias-Martinez. Socio-economic levels and human mobility, Qual Meets Quant Workshop – QMQ 2010 at the int.conf on information & communication technologies and development, ICTD (2010).
- [16] Rubio, A., Frias-Martinez, V., Frias-Martinez E. & Oliver, N. (2010, March). Human Mobility in Advanced and Developing Economies: A Comparative Analysis, in AAAI Spring Symposia Artificial Intelligence for Development, AID, Stanford, USA.
- [17] Eagle, N., Macy, M. & Claxton, R. (2010). Network Diversity and Economic Development, Science 328 2029-1030.



# **Paper 3**

---



# INTERFACE

rsif.royalsocietypublishing.org

## Research



**Cite this article:** Steele JE *et al.* 2017

Mapping poverty using mobile phone and satellite data. *J. R. Soc. Interface* **14**: 20160690.  
<http://dx.doi.org/10.1098/rsif.2016.0690>

Received: 29 August 2016

Accepted: 3 January 2017

### Subject Category:

Life Sciences – Mathematics interface

### Subject Areas:

biometrics, environmental science,  
biomathematics

### Keywords:

poverty mapping, mobile phone data, Bayesian  
geostatistical modelling, remote sensing

### Author for correspondence:

Jessica E. Steele

e-mail: [steele.jessica@gmail.com](mailto:steele.jessica@gmail.com)

Electronic supplementary material is available  
online at <https://dx.doi.org/10.6084/m9.fig-share.c.3662800>.

**THE ROYAL SOCIETY**  
PUBLISHING

# Mapping poverty using mobile phone and satellite data

Jessica E. Steele<sup>1,2</sup>, Pål Roe Sundsøy<sup>3</sup>, Carla Pezzulo<sup>1</sup>, Victor A. Alegana<sup>1</sup>,  
Tomas J. Bird<sup>1</sup>, Joshua Blumenstock<sup>4</sup>, Johannes Bjelland<sup>3</sup>,  
Kenth Engø-Monsen<sup>3</sup>, Yves-Alexandre de Montjoye<sup>5</sup>, Asif M. Iqbal<sup>6</sup>,  
Khandakar N. Hadiuzzaman<sup>6</sup>, Xin Lu<sup>2,7,8</sup>, Erik Wetter<sup>2,9</sup>, Andrew J. Tatem<sup>1,2,10</sup>  
and Linus Bengtsson<sup>2,7</sup>

<sup>1</sup>Geography and Environment, University of Southampton, University Road, Building 44, Southampton, UK

<sup>2</sup>Flowminder Foundation, Roslagsgatan 17, Stockholm, Sweden

<sup>3</sup>Telenor Group Research, Oslo, Norway

<sup>4</sup>School of Information, University of California, Berkeley, CA, USA

<sup>5</sup>Data Science Institute, Imperial College London, London, UK

<sup>6</sup>Grameenphone Ltd, Dhaka, Bangladesh

<sup>7</sup>Public Health Sciences, Karolinska Institute, Stockholm, Sweden

<sup>8</sup>College of Information System and Management, National University of Defense Technology,  
Changsha, Hunan, People's Republic of China

<sup>9</sup>Stockholm School of Economics, Saltmätargatan 13-17, Stockholm, Sweden

<sup>10</sup>John E Fogarty International Center, National Institutes of Health, Bethesda, MD, USA

JES, 0000-0001-6741-1195; PRS, 0000-0003-4596-3209; VAA, 0000-0001-5177-9227;  
TJB, 0000-0003-1345-6480; JB, 0000-0002-1813-7414; KE-M, 0000-0003-1618-7597;  
Y-AdM, 0000-0002-2559-5616; LB, 0000-0001-9712-2669

Poverty is one of the most important determinants of adverse health outcomes globally, a major cause of societal instability and one of the largest causes of lost human potential. Traditional approaches to measuring and targeting poverty rely heavily on census data, which in most low- and middle-income countries (LMICs) are unavailable or out-of-date. Alternate measures are needed to complement and update estimates between censuses. This study demonstrates how public and private data sources that are commonly available for LMICs can be used to provide novel insight into the spatial distribution of poverty. We evaluate the relative value of modelling three traditional poverty measures using aggregate data from mobile operators and widely available geospatial data. Taken together, models combining these data sources provide the best predictive power (highest  $r^2 = 0.78$ ) and lowest error, but generally models employing mobile data only yield comparable results, offering the potential to measure poverty more frequently and at finer granularity. Stratifying models into urban and rural areas highlights the advantage of using mobile data in urban areas and different data in different contexts. The findings indicate the possibility to estimate and continually monitor poverty rates at high spatial resolution in countries with limited capacity to support traditional methods of data collection.

## 1. Background

In 2015, approximately 700 million people lived in extreme poverty [1]. Poverty is a major determinant of adverse health outcomes including child mortality [2], and contributes to population growth [3], societal instability and conflict [4]. Eradicating poverty in all its forms remains a major challenge and the first target of the Sustainable Development Goals (SDGs) [5]. To eradicate poverty, it is crucial that information is available on where affected people live. Such data improve the understanding of the causes of poverty, enable improved allocation of resources for poverty alleviation programmes, and are a critical

component for monitoring poverty rates over time. The latter issue is especially pertinent for efforts aimed at reaching the SDGs, which need to be monitored at national and subnational levels over the coming 15 years [5].

The definition of poverty and the measurement methods used to identify poor persons are part of a longstanding discussion in development economics [6–9]. Different approaches exist to calculate indicators of living standards, including the construction of unidimensional and multidimensional indices, as well as the use of monetary or non-monetary metrics. A further discussion for living standard indices regards the methods used to set appropriate thresholds (poverty lines) under which a person is defined as poor [10–12]. Monetary-based metrics identify poverty as a shortfall in consumption (or income) and measure whether households or individuals fall above or below a defined poverty line [13,14]. By contrast, asset-based indicators define household welfare based on asset ownership (e.g. refrigerator, radio or bicycle), dwelling characteristics, and access to basic services like clean water and electricity [15]. Moreover, poverty indicators can capture the status of a household or individual at a given point in time, or identify chronic versus transient poverty over time [14,16–18].

Every approach used to calculate indicators of living standards for a population has its advantages and disadvantages, and each indicator discerns different characteristics of the population. Consumption data can be highly noisy due to recall error or because expenditures occurred outside the period captured in surveys, but provide a better shorter-term concept of poverty [19,20]. Asset-based measures have been regarded as a better proxy for the long-term status of households as they are thought to be more representative of permanent income or long-term control of resources [20–22]. The same population can be ranked quite differently along a poverty distribution when comparing consumption and asset-based measures and many assumptions are necessarily accepted in order to do such comparisons. These include assumptions that the data represent the same populations in the same time period; that the indicators are well matched in their wording and response options; and that the poverty measures have a similar distribution of responses [20,23]. Furthermore, it is difficult to compare asset-based measures to income or consumption as it is not straightforward to link the productive potential of a household to their assets owned; this can be particularly relevant in rural areas where the return on physical assets can be strongly environmentally related and interactions among assets may be important [24]. These factors necessitate a flexible approach to modelling poverty as indicators representing asset-based, consumption-based and income-based measures are not necessarily expected to produce similar results.

While numerous high-resolution indicators of human welfare are routinely collected for populations in high-income countries, the geographical distribution of poverty in low- and middle-income countries (LMICs) is often uncertain [25]. Small area estimation (SAE) forms the standard approach to produce sub-national estimates of the proportion of households in poverty. SAE uses statistical techniques to estimate parameters for sub-populations by combining household survey and census data to use the detail in household surveys and the coverage of the census. Common variables between the two are used to predict a poverty metric across the population [26–28]. These techniques rely on the availability of census data, which are typically collected every 10 years and often

released with a delay of one or more years, making the updating of poverty estimates challenging. Recently, there are promising signs that novel sources of high-resolution data can provide an accurate and up-to-date indication of living conditions. In particular, recent work illustrates the potential of features derived from remote sensing and geographic information system data [29–35] (hereafter called RS data) and mobile operator call detail records (CDRs) [36–39]. However, the predictive power in integrating these two data sources, and their ability to estimate different measures of poverty has not been evaluated.

RS and CDR data capture distinct and complimentary correlates of human living conditions and behaviour. For example, RS data of physical properties, such as rainfall, temperature and vegetation capture information related to agricultural productivity, while distance to roads and cities reflects access to markets and information. Similarly, monthly credit consumption on mobile phones and the proportion of people in an area using mobile phones indicate household access to financial resources, while movements of mobile phones and the structure and geographical reach of the calling networks of individuals may be correlated with remittance flows and economic opportunities [39–41].

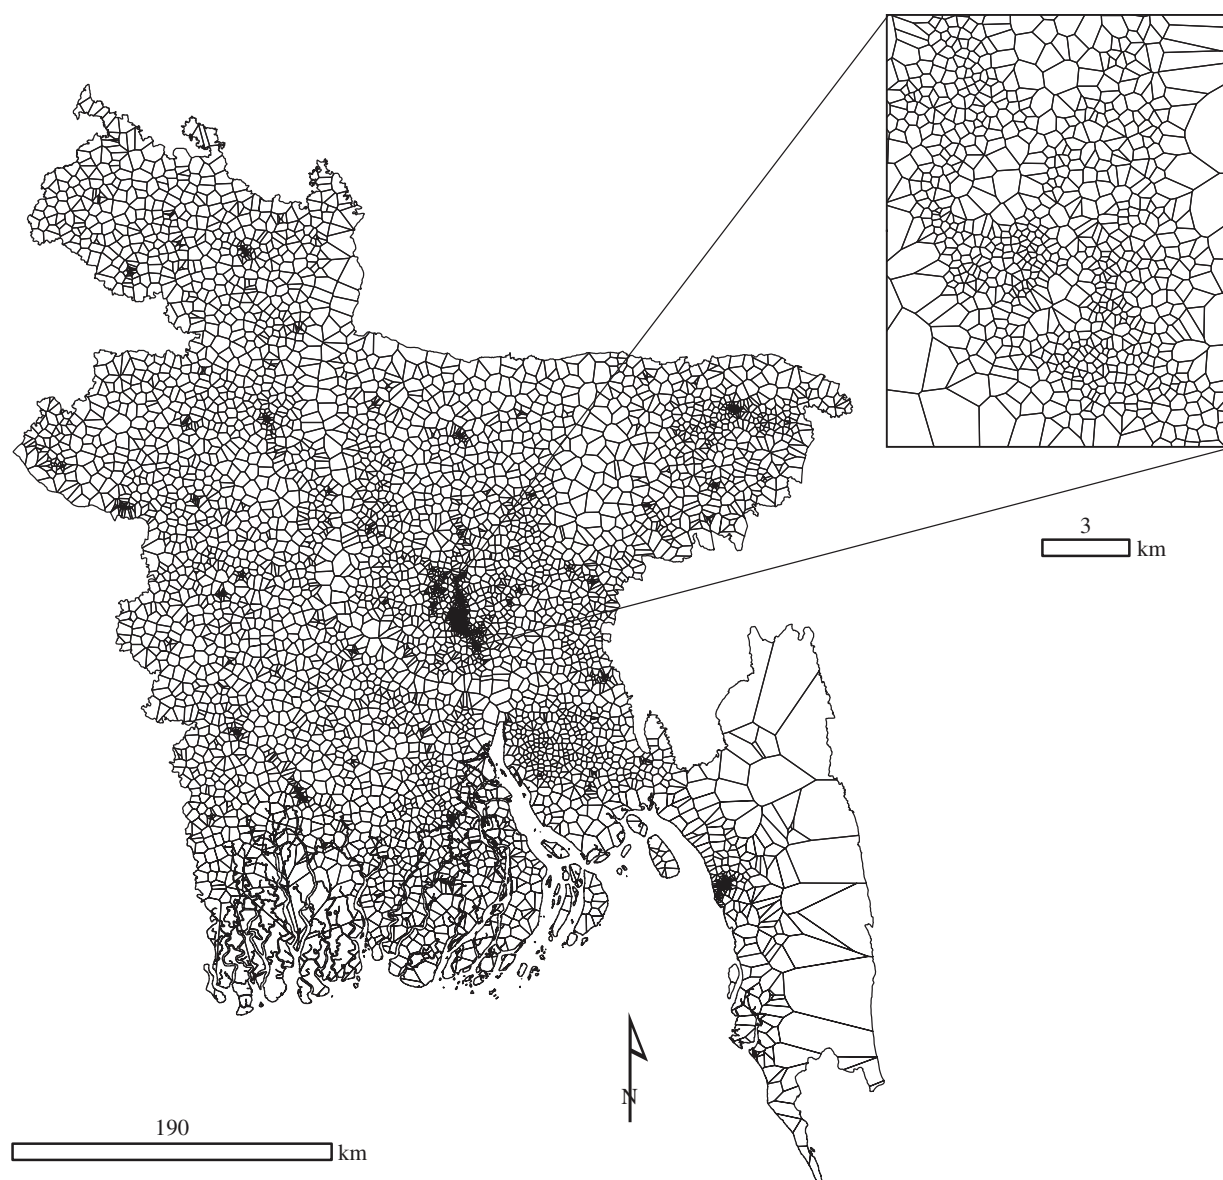
RS and CDR data are generated at different spatial scales, which further complement each other. The CDR indicators used in this study are derived from data aggregated at the level of the physical cell towers to preserve the privacy of individual subscribers. Thus, the spatial resolution of these data is determined by tower coverage, which is larger in rural areas and fine-scaled in urban areas. By contrast, RS data can be relatively coarse in urban areas and only capture physical properties of the land. As RS and CDR data are continually collected, the ability to produce accurate maps using these data types offers the promise of ongoing subnational monitoring required by the SDGs.

Here, we use overlapping sources of RS, CDR and traditional survey-based data from Bangladesh to provide the first systematic evaluation of the extent to which different sources of input data can accurately estimate three different measures of poverty. To date, the predictive power in integrating these data sources, and their ability to estimate different measures of poverty, has not been evaluated. We use hierarchical Bayesian geostatistical models (BGMs) to construct highly granular maps of poverty for three commonly used indicators of living standards: the Demographic and Health Surveys (DHS) Wealth Index (WI); an indicator of household expenditures (Progress out of Poverty Index, PPI) [42] and reported household monetary income. We additionally compare our results with previous poverty estimates for Bangladesh at coarser and finer resolutions.

## 2. Material and methods

### 2.1. Spatial scale and data processing

All data used in this study were processed to ensure that projections, resolutions and extents matched. The spatial scale of analysis was based on approximating the mobile tower coverage areas using Voronoi tessellation [43] and models were built on the scale of the Voronoi polygons (figure 1). This allowed us to maintain the fine spatial detail in mobile phone data within urban areas, as Voronoi polygon size, and corresponding spatial detail, varies greatly from urban to rural areas (minimum 60 m,



**Figure 1.** Spatial structure of Voronoi polygons based on the configuration of mobile phone towers in Bangladesh. The zoom window shows the spatial detail of Dhaka.

maximum 5 km) as shown in the figure. All datasets were then summarized to spatially align with these polygons. In practice, each polygon was assigned RS and CDR values representing the mean, sum or mode of the corresponding data. The survey data are matched to the Voronoi based on the GPS located lat/long of PPI data, the lat/long representing the centroid of each DHS cluster, and the home tower of each income survey respondent. Where multiple points from the same output (WI, PPI and income) fell within the same polygon, we used the mean aggregated value.

## 2.2. Poverty data

We used three geographically referenced datasets representing asset, consumption and income-based measures of wellbeing in Bangladesh (see the electronic supplementary material, figure S1 and section A.1). These data were obtained from three sources: the 2011 Bangladesh DHS, the 2014 FII survey [44] with data collected on the PPI ([www.progressoutofpoverty.org](http://www.progressoutofpoverty.org)) and national household surveys conducted by Telenor Group subsidiary Grameenphone (GP) between November 2013 and March 2014 collecting household income data.

The DHS WI is constructed by taking the first principal component of a basket of household assets and housing characteristics such as floor type and ceiling material, which explains the largest percentage of the total variance, adjusting for differences in urban and rural strata [45]. A final composite combined score is then used as a WI whereby each household is assigned its correspondent quintile in the distribution and each individual belonging to the same household shares the same WI score. A higher score implies higher socioeconomic status (range =  $-1.45$  to  $3.5$ ). Here, we used aggregated average WI scores per primary sampling unit (PSU) for 600 PSUs (207 in urban areas and 393 in rural areas) to estimate the mean WI of sampled populations residing in each Voronoi polygon.

The PPI is a measurement tool built from the answers to 10 questions about a household's characteristics and asset ownership, scored to compute the likelihood the household is living above or below a poverty line. In Bangladesh, these poverty scorecard questions were determined using data from the 2010 Household Income and Expenditure Survey (HIES) [42,46], and used in a nationally representative survey of 6000 Bangladeshi adults undertaken in 2014 [44]. Together with basic demographics and access to financial services information, the 10 questions



needed to construct the PPI were collected. These data were used to assign a poverty measure to each individual interviewed: the likelihood they have *per capita* expenditure above or below a poverty line. Here, we estimate the mean likelihood (range = 12.3–99.7%) of populations residing in each Voronoi polygon to be below the \$2.50 a day poverty line.

Income data were obtained from two independent, sequential household surveys run by GP. For each survey, face-to-face interviews were conducted with 90 000 individuals, and their corresponding household income was collected, together with basic demographic information for each survey participant (e.g. gender, age, profession, education) and phone usage. Respondents were directly asked about income and were requested to place themselves within pre-set income bins. Among GP subscribers, CDRs were successfully linked to phone numbers for 76 000 participants. Here we converted income bins to USD (range = 0–1285\$) and modelled the average USD for each Voronoi polygon.

### 2.3. CDR and RS data

CDR features were generated from four months of mobile phone metadata collected between November 2013 and March 2014. GP subscribers consented to the use of their data for the analysis. GP, the largest mobile network operator in Bangladesh, had 48 million customers at the time of the analysis, with a network covering 99% of the population and 90% of the land area [47]. CDR features range from metrics such as basic phone usage, top-up patterns, and social network to metrics of user mobility and handset usage. These features are easily made available in data warehouses and do not rely on complex algorithms. They include various parameters of the corresponding distributions such as weekly or monthly median, mean and variance (see the electronic supplementary material).

We further identified, assembled and processed 25 raster and vector datasets into a set of RS covariates for the whole of Bangladesh at a 1 km spatial resolution. These data were obtained from existing sources and produced *ad hoc* for this study to include environmental and physical metrics likely to be associated with human welfare [31,33,48–50] such as vegetation indices, night-time lights, climatic conditions, and distance to roads or major urban areas. A full summary of assembled covariates is provided in the electronic supplementary material.

### 2.4. Covariate selection

Prior to statistical analyses, all CDR and RS covariate data were log transformed for normality. Bivariate Pearson's correlations were computed for each pair of covariates to assess multicollinearity, and for high correlations ( $r > 0.70$ ), we eliminated covariates that were less generalizable outside Bangladesh. For example, population data are widely available (e.g. [www.worldpop.org.uk/](http://www.worldpop.org.uk/)) but births data may not be; similarly, volumes of calls could be computed and compared across countries, but charges may be country-specific.

To identify the set of predictors most suitable for modelling the WI, PPI, and income data, we employed a model selection stage as is common in statistical modelling [51]. For this we used non-spatial generalized linear models (glms), implemented via the R *glmulti* package [52,53], to build every possible non-redundant model for every combination of covariates. Models were built on a randomly selected 80% of the data to guard against overfitting. Models were chosen using Akaike's information criterion (AIC), which ranks models based on goodness of fit and complexity, while penalizing deviance [52]. A full IC-based approach such as this allows for multi-model inference. Where multiple glms had near-identical AIC values, we selected the model with the fewest number of covariates. For the CDR data only, we used forward and backward stepwise selection ( $p = 0.05$ )

prior to model selection to reduce the initial CDR inputs from 150 to 30 or less. The covariate selection process was completed for all three poverty measures for national, urban and rural strata, and using RS-only, CDR-only and CDR–RS datasets (27 resulting models). This allowed us to explore differences in factors related to urban and rural poverty, as well as to explicitly compare the ability of RS-only, CDR-only and CDR–RS datasets to predict poverty measures. The resulting models were then used in the hierarchical Bayesian geostatistical approach (see the electronic supplementary material, tables S2a–c).

### 2.5. Prediction mapping

Using the models selected by the previous step, we employed hierarchical Bayesian geostatistical models (BGMs) to predict the three poverty metrics at unsampled locations across the population. We chose BGMs as they offer several advantages for addressing the limitations and constraints associated with modelling geolocated survey data. These include straightforwardly imputing missing data, allowing for the specification of prior distributions in model parameters and spatial covariance, and estimating uncertainty in the predictions as a distribution around each estimate [54,55].

Additionally, we needed to account for spatial autocorrelation in the data as they are aligned to the tower locations, which are clustered across varying spatial scales (described in §2.1 and figure 1). BGMs can achieve this through incorporating a spatially varying random effect. Here, the Voronoi polygons themselves form the neighbourhood structure for this spatial random effect, and neighbours are defined within a scaled precision matrix [56]. The matrix represents the spatially explicit processes that may affect poverty estimates. It is passed through a graph function in the model which assumes the neighbour relations are connected [57], that is all adjacent polygons share a boundary. This function accounts for the spatial covariance in the data by allowing observations to have decreasing effects on predictions that are further away.

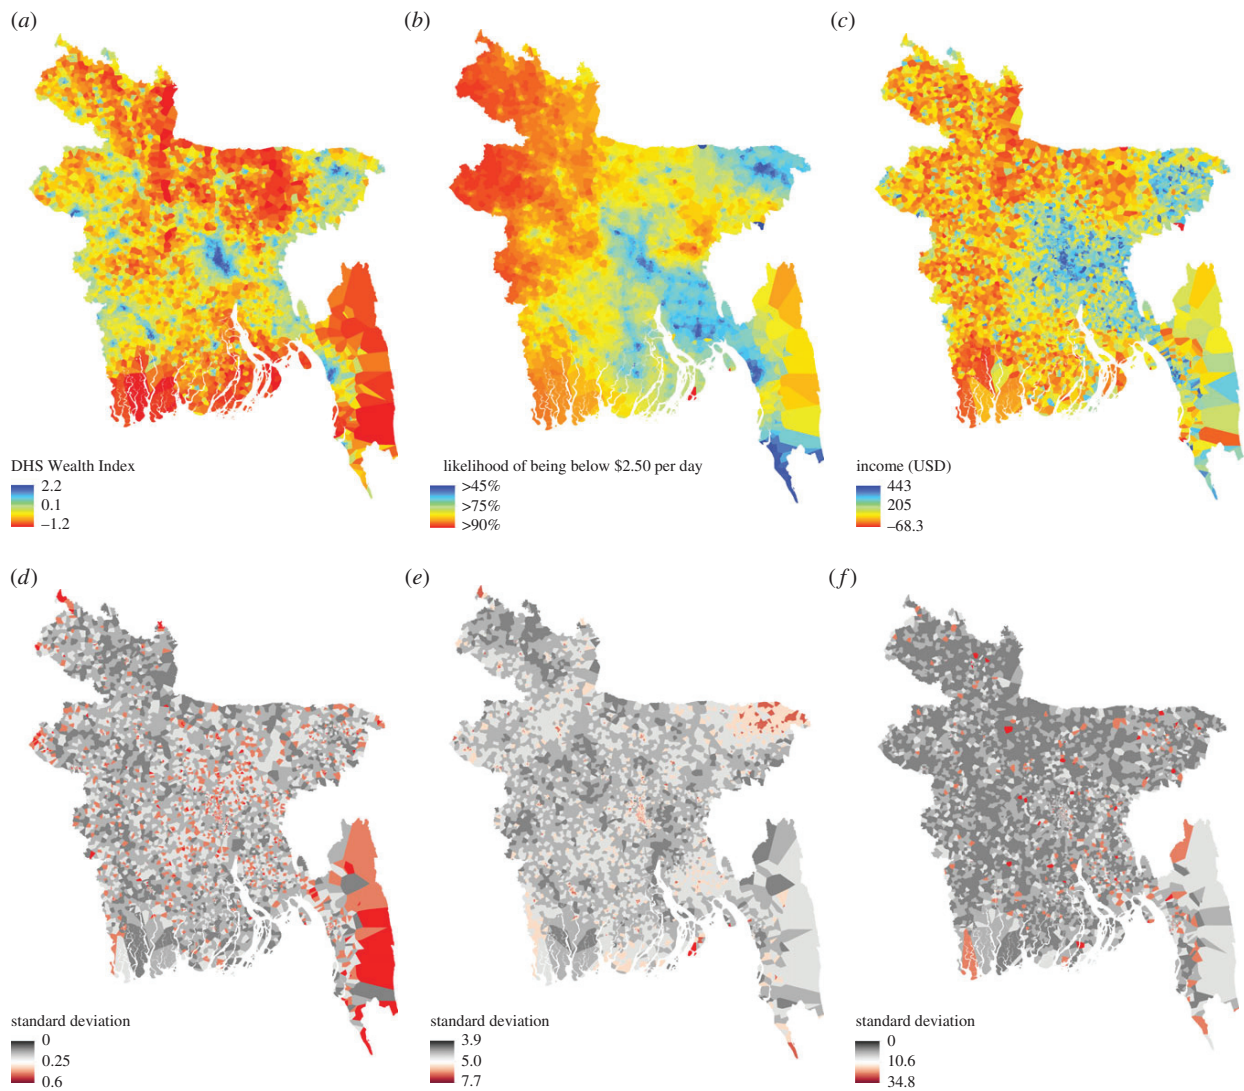
Here, all BGMs were implemented using integrated nested Laplace approximations (INLA) [58], which uses an approximation for inference and avoids the computational demands, convergence issues and mixing problems sometimes encountered by MCMC algorithms [59]. The model is fit using R-INLA, with the Besag model for spatial effects specified inside the function [60,61]. In the Besag model, Gaussian Markov random fields (GMRFs) are used as priors to model spatial dependency structures and unobserved effects. GMRFs penalize local deviation from a constant level based on the precision parameter  $\tau$ , where the hyperpriors are loggamma distributed [56]. The hyperprior distribution governs the smoothness of the field used to estimate spatial autocorrelation [56]. The spatial random vector  $\mathbf{x} = (x_1, \dots, x_n)$  is thus defined as

$$x_i | x_{i \sim j}, \tau \sim \mathcal{N} \left( \frac{1}{n_i} \sum_{i \sim j} x_j, \frac{1}{n_i \tau} \right),$$

where  $n_i$  is the number of neighbours of node  $i$ ,  $i \sim j$  indicates that the two nodes  $i$  and  $j$  are neighbours. The precision parameter  $\tau$  is represented as

$$\theta_1 = \log \tau,$$

where the prior is defined on  $\theta_1$  [60]. The geostatistical models defined for the WI, PPI and income data were applied to produce predictions of the each poverty metric for each Voronoi polygon as a posterior distribution with complete modelled uncertainty around estimates. The posterior mean and standard deviation for each polygon were then used to generate prediction maps with associated uncertainty (figure 2 and electronic supplementary material, figures S2–S6). Model performance was based on out-of-sample validation statistics calculated on a 20% test subset of data. Pearson product-moment correlation coefficient ( $r$ ) (or



**Figure 2.** National level prediction maps for mean WI (a) with uncertainty (d); mean probability of households being below \$2.50/day (b) with uncertainty (e); and mean USD income (c) with uncertainty (f). Maps were generated using call detail record features, remote sensing data and Bayesian geostatistical models. The maps show the posterior mean and standard deviation from CDR–RS models for the WI and income data (a,c), and the RS model for the PPI (b). Red indicates poorer areas in prediction maps, and higher error in uncertainty maps.

Spearman's rho ( $\rho$ ) for  $n < 100$ ), root-mean-square-error (RMSE), mean absolute error (MAE) and the coefficient of determination ( $r^2$ ) were calculated for all BGMs. Finally, because glms do not incorporate prior information for model parameters, we ran each model through INLA while excluding the random spatial effect to obtain non-spatial Bayesian estimates and compare model fit and performance due to the explicit spatial process.

### 3. Results

We find models employing a combination of CDR and RS data generally provide an advantage over models based on either data source alone. However, RS-only and some CDR-only models performed nearly as well (table 1). While the combined CDR–RS model performed well in both urban ( $r^2 = 0.78$ ) and rural ( $r^2 = 0.66$ ) areas, and at the national level ( $r^2 = 0.76$ ), the performance of RS-only and CDR-only models was more context-dependent. For example, PPI and income models did not improve predictions in urban areas, but in rural areas the RS-only models performed nearly as well for both indicators. The fine spatial granularity of the resultant poverty

estimates can be shown in figure 2, which shows the predicted distribution of poverty for all three measures. Spatially, the models exhibit higher uncertainty where fewer data are available, such as the peninsular areas surrounding Chittagong in the southeast where mobile towers are sparse. We also find that explicitly modelling the spatial covariance in the data was critically important. This resulted in improved predictions, lower error and better measures of fit based on cross-validation and the deviance information criteria (DIC), a hierarchical modelling generalization of the AIC [62] (electronic supplementary material, tables S3 and S4).

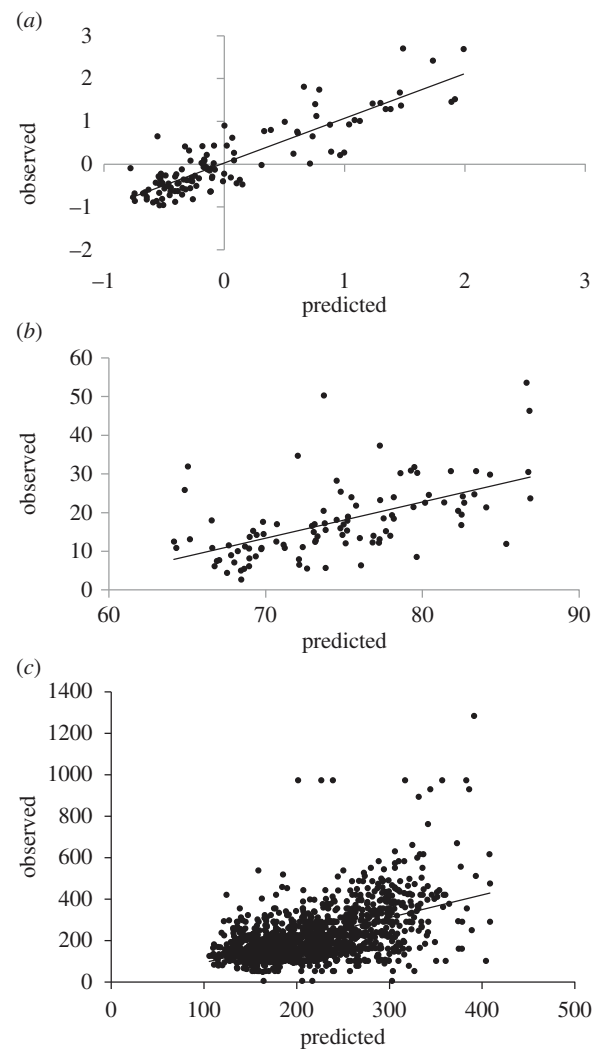
Separating estimation by urban and rural regions further highlights the importance of different data in different contexts (electronic supplementary material, tables S2a–c). Night-time lights, transport time to the closest urban settlement, and elevation were important nationally and in rural models; climate variables were also important in rural areas. Distances to roads and waterways were significant in urban and rural strata. In general, the addition of CDR data did not change the selection of RS covariates at any level. Top-up features derived from recharge amounts and tower averages

**Table 1.** Cross-validation statistics based on a random 20% test subset of data for all Bayesian geostatistical models.

poverty metric	model	$r^2$	RMSE
whole country			
DHS WI	CDR-RS	0.76	0.394
	CDR	0.64	0.483
	RS	0.74	0.413
PPI	CDR-RS	0.25	57.907
	CDR	0.23	58.562
	RS	0.32	57.439
income	CDR-RS	0.27	105.465
	CDR	0.24	107.155
	RS	0.22	108.682
urban			
DHS WI	CDR-RS	0.78	0.424
	CDR	0.70	0.552
	RS	0.71	0.433
PPI	CDR-RS	0.00	60.128
	CDR	0.03	60.935
	RS	0.00	60.384
income	CDR-RS	0.15	168.452
	CDR	0.15	172.738
	RS	0.05	176.705
rural			
DHS WI	CDR-RS	0.66	0.402
	CDR	0.50	0.483
	RS	0.62	0.427
PPI	CDR-RS	0.18	57.397
	CDR	0.17	57.991
	RS	0.21	57.162
income	CDR-RS	0.14	81.979
	CDR	0.13	82.773
	RS	0.23	76.527

were significant in every model, affirming their importance in poverty work. People predicted to be poorer top-up their phones more frequently in small amounts. Per cent nocturnal calls, and count and duration of SMS traffic were significant nationally. Mobility and social network features were important in all three strata. In urban areas, SMS traffic was important, whereas multimedia messaging and video attributes were key in rural areas.

Models were most successful at reconstructing the WI to model poverty ( $r^2 = 0.76$ ); consumption-based and income-based poverty proved more elusive. WI models have better fit, lower error and higher explained variance based on out-of-sample validation (figure 3). Combined CDR-RS data produced the best WI models and lowest error ( $r^2_{\text{(CDR-RS)}} = 0.76$ ,  $r^2_{\text{(RS)}} = 0.74$ ,  $r^2_{\text{(CDR)}} = 0.64$ ;  $\text{RMSE}_{\text{(CDR-RS)}} = 0.394$ ,  $\text{RMSE}_{\text{(RS)}} = 0.413$ ,  $\text{RMSE}_{\text{(CDR)}} = 0.483$ ). However, for the PPI models, the best model predicting the probability of falling below \$2.50/day was the RS-only model (figure 2*b,c*,  $r^2_{\text{(RS)}} = 0.32$ ;  $\text{RMSE}_{\text{(RS)}} = 57.439$ ). The

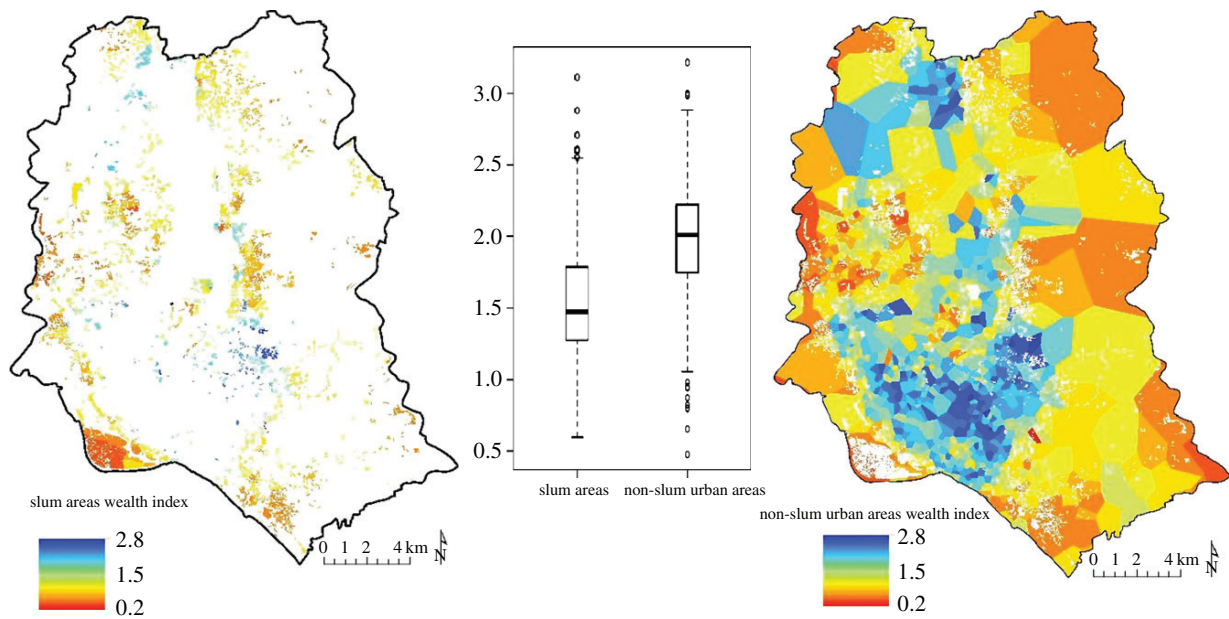


**Figure 3.** Out-of-sample observed versus predicted values for (a) DHS WI using mobile phone and remote sensing data:  $r^2 = 0.76$ ,  $n = 117$ ,  $p < 0.001$ ,  $\text{RMSE} = 0.394$ ; (b) progress out of Poverty Index using remote sensing data:  $r^2 = 0.32$ ,  $n = 100$ ,  $p < 0.001$ ,  $\text{RMSE} = 57.439$ ; and (c) income using mobile phone and remote sensing data:  $r^2 = 0.27$ ,  $n = 1384$ ,  $p < 0.001$ ,  $\text{RMSE} = 105.465$ .

model discerns many urban areas but also predicts areas with very low poverty likelihood and high uncertainty outside urban areas, especially around Sylhet in the northeast. Income predictions (figure 2*c,f*) show greater variation across the country, and the best national model was for combined CDR-RS data ( $r^2_{\text{(CDR-RS)}} = 0.27$ ,  $\text{RMSE}_{\text{(CDR-RS)}} = 105.465$ ).

The resulting predictions line up well with existing SAE estimates for Bangladesh, and with high-resolution maps of slum areas in Dhaka. The urban CDR-RS model has the highest explained variance for any model ( $r^2_{\text{(CDR-RS-urb)}} = 0.78$ ) and the urban CDR-only model outperforms the national CDR-only model ( $r^2_{\text{(CDR-urb)}} = 0.70$ ). Precision and accuracy are slightly lower, but the improved correlation highlights the advantage of using CDRs within a diverse urban population. To explore this further, we compared our WI predictions against a spatially explicit dataset of slum areas in Dhaka [63] (figure 4). We find the mean predicted WI of slum and non-slum areas to be significantly different,  $t_{615} = -17.2$ ,  $p < 0.001$ , predicting slum areas to be poorer than non-slum areas.





**Figure 4.** Comparison of predicted mean DHS WI values between slum and non-slum areas in Dhaka as delineated by Gruebner *et al.* [63]  $t_{615} = -17.2$ ,  $p < 0.001$ . The 95% confidence interval using Student's  $t$ -distribution with 615 degrees of freedom is  $(-0.48, -0.38)$ .

To compare our method to previous poverty estimates at administrative level 3 (upazila), we used the same methodology at the lower spatial resolution, using the upazila boundaries to form the random spatial effect in the model, and covariates from the best national level model for each poverty measure. We find strong correlations ( $r = -0.91$  and  $-0.86$  for the WI;  $0.99$  and  $0.97$  for the PPI; and  $-0.96$  and  $-0.94$  for income, respectively,  $p < 0.001$  for all models) between our upazila predictions and earlier estimates of poverty derived from SAE techniques based on data from the 2010 Household Income and Expenditure (HIES) survey and 2011 census [64] (figure 5). The  $r$ -values reported for WI and income are negative at administrative level 3 because as the proportion of people below the poverty line as estimated by Ahmed *et al.* decreases, the WI value and income in USD of the sampled population increases. That is, people who are wealthier as estimated by the WI and income data are also less likely to live below the poverty line according to earlier estimates. The geostatistical method presented here thus accurately maps heterogeneities at small spatial scales while correlating well with earlier coarser estimates. All remaining WI, PPI and income prediction maps are provided in the electronic supplementary material.

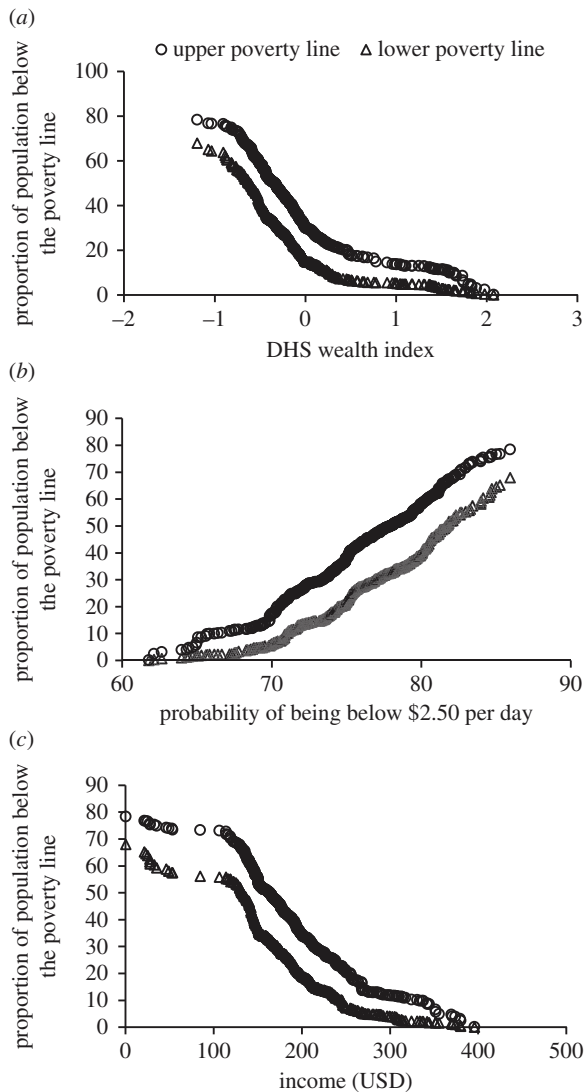
## 4. Discussion

This work represents the first attempt to build predictive maps of poverty using a combination of CDR and RS data. The results demonstrate that CDR-only and RS-only models perform comparably in their ability to map poverty indicators, and that integrating these data sources provides improvement in predictive power and lower error. These results are promising as the CDR data here produce accurate, high-resolution estimates in urban areas not possible using RS data alone. As such, CDRs potentially allow for estimation of wealth at much finer granularity—including the neighbourhood or even the household or individual—than the current generation

of RS technologies [36]. While CDRs are proprietary data, they are increasingly used in research, and have formed the basis for hundreds of published articles over the past few years [65]. They also provide significant advantages in temporal granularity: CDRs update in real-time versus RS data, which update far less frequently. Although in this study we have not used dynamic validation data, it is a clear future application for CDRs in real-time to better comprehend the dynamic nature of poverty.

The higher accuracy of predictions for the asset-based WI over other poverty metrics is presumably due to several factors. The predictive power for assets has been shown to be higher than for consumption [35] in addition to the aforementioned issues of survey question wording and response options [20,23]. Further, income and consumption can vary hugely by day, week, and can be related to changes in household size, job loss or gain, piecework or harvest outcomes. Assets and housing characteristics are generally considered more stable [20–22]. For the datasets used in this study, WI data are based on clusters of households, and this sampling strategy provides more robust estimates and less variability than the individually based PPI and income data. Greater success in predicting the WI is also presumably due to the WI measuring a wider range of living standard across the population. That is, the full range of distribution from poorest to wealthiest in the population is represented in these data. Alternatively, by considering a streamlined 10 questions, the PPI is meant to identify the poorest individuals in a population. Similarly, in the income data, there were very few respondents in higher income categories.

The higher error associated with CDR-only models is not surprising considering the noise inherent in these data. CDR features are derived from daily and weekly measurements aggregated over short temporal intervals, while RS covariates are generally comprised of long-term averages or comparatively less dynamic measures of location and access such as roads or proximity to urban centres. Bearing this in mind,



**Figure 5.** Comparison of the proportion of people falling below upper (circles) and lower (triangles) poverty lines estimated by Ahmad *et al.* [64] and (a) predicted mean WI using mobile phone and remote sensing data, (b) predicted probability of being below \$2.50 per day using remote sensing data and (c) predicted income using mobile phone and remote sensing data. All models were predicted at the upazila scale (Admin unit 3). Pearson's  $r$  correlations:  $-0.91$  and  $-0.86$  for the WI;  $0.99$  and  $0.97$  for the PPI; and  $-0.96$  and  $-0.94$  for income, respectively ( $p < 0.001$  for all models).

we find CDR data useful for estimating poverty in the absence of ancillary datasets.

Our findings provide further support for correlations between socio-economic measures and night-time light intensity [36,48,49], access to roads and cities [50,66], entropy of contacts [37,40] and mobility features [39]. The universal coverage of cell towers across Bangladesh made it possible to predict poverty at high-resolution in both urban and rural areas. Within urban areas, the high correlation with maps of slums in Dhaka suggests we are capturing the poorest populations. Even if the poorest populations are not generating call data [36], and thus not included in the CDRs, we still see a clear difference in WI predictions between slum and non-slum areas using tower level CDR aggregates. This finding extends recent work which predicted wealth and poverty at the district level, but were unable to verify predictions at finer scales [36].

Using CDRs and RS data within BGMs to produce accurate, high-resolution poverty maps in LMICs offers a way to complement census-based methods and provide more regular updates. Regularly updated poverty estimates are necessary to enable subnational monitoring of the SDGs during inter-censal years and are critical to ensure mobilization of resources to end poverty in all its dimensions as set out in SDG 1. Poverty estimates are time sensitive and become obsolete when factors such as migration rates, infrastructure, and market integration change [67]. Furthermore, the methods presented here offer a workaround to estimating poverty with household survey data, which can be time consuming and expensive to obtain.

To end poverty in all its dimensions, it is likely that methods that exploit information from, and correlations between, many different data sources will provide the greatest benefit in understanding the distribution of human living conditions. To leverage data from differing sample sizes, temporal and spatial scales, BGMs provide such a rigorous framework. This study further provides an example of how aggregated CDR data can be processed in such a way that detailed maps can be created without revealing sensitive user and commercial information. As insights from CDRs and other remote sensing data become more widely available, analysing these data at regular intervals could allow for dynamic poverty mapping and provide the means for operationally monitoring poverty. The combination of spatial detail and frequent, repeated measurements may distinguish the transitorily poor from the chronically poor, and allow for monitoring economic shocks [68]. This offers the potential for a fuller characterization of the spatial distribution of poverty and provides the foundation for evidence-based strategies to eradicate poverty. Researchers would do well to use the additional information and granularity afforded by CDR data with matched individual-based consumption data to further infer novel and useful information from mobile data.

**Authors' contributions.** J.E.S. was responsible for research design, production of RS covariates, data cleaning and processing, statistical analyses, interpretation, drafting and production of the final manuscript. C.P. was responsible for survey data management, cleaning and processing, and interpretation and drafting of the final manuscript. P.S., J.Bj. and K.E. were responsible for CDR data management, cleaning, and production of CDR data, and interpretation and drafting of the final manuscript. J.Bl. was responsible for interpretation, drafting, and production of the final manuscript. V.A., T.B., Y.M., X.L. and E.W. were responsible for interpretation and production of the final manuscript. A.I. and K.N.H. were responsible for income survey data collection, management and processing. A.J.T. and L.B. were responsible for overall scientific management, interpretation and production of the final manuscript. All authors gave final approval for publication.

**Competing interests.** We declare we have no competing interests.

**Funding.** J.E.S., E.W. and J.Bl. are supported by the Bill & Melinda Gates Foundation (OPP1106936). C.P. is supported by the Bill & Melinda Gates Foundation (OPP1106427). X.L. acknowledges the Natural Science Foundation of China under grant nos. 71301165 and 71522014. A.J.T. is supported by funding from NIH/NIAID (U19AI089674), the Bill & Melinda Gates Foundation (OPP1106427, 1032350, OPP1134076 and OPP1094793), the Clinton Health Access Initiative, National Institutes of Health, and a Wellcome Trust Sustaining Health grant (106866/Z/15/Z). L.B. acknowledges the Swedish Research Council, grant no. D0313701. This work forms part of the WorldPop Project ([www.worldpop.org.uk](http://www.worldpop.org.uk)) and Flowminder Foundation ([www.flowminder.org](http://www.flowminder.org)).

**Acknowledgements.** The authors gratefully acknowledge Dr Elisabeth zu Erbach-Schoenberg, Dr Nick Ruktanonchai and Dr Alessandro Soricetta for useful discussions. We also wish to thank two anonymous reviewers for providing useful comments. The funders had no role in study design, data collection and analysis, decision to publish or preparation of the manuscript.

## References

- Cruz M, Foster J, Quillin B, Schellekens P. 2015 *Ending extreme poverty and sharing prosperity: progress and policies*. 83. Washington, DC: World Bank Development Economics Group.
- Målvist M. 2015 Abolishing inequity, a necessity for poverty reduction and the realisation of child mortality targets. *Arch. Dis. Child.* **100**, S5–S9. (doi:10.1136/archdischild-2013-305722)
- Population and poverty | UNFPA - United Nations Population Fund. 2016 See <http://www.unfpa.org/resources/population-and-poverty> (accessed: 21 January 2016).
- Braithwaite A, Dasandi N, Hudson D. 2014 Does poverty cause conflict? Isolating the causal origins of the conflict trap. *Confl. Manag. Peace Sci.* **33**, 45–66. (doi:10.1177/0738894214559673)
- United Nations General Assembly. 2015 Transforming our world: the 2030 Agenda for Sustainable Development.
- Chambers R. 2006 What is poverty? Who asks? Who answers? Poverty in Focus December 2006, UNDP International Poverty Centre. See <http://opendocs.ids.ac.uk/opendocs/handle/123456789/120>.
- Deaton A, Zaidi S. 1999 *Guidelines for constructing consumption aggregates for welfare analysis*. Princeton, NJ: Woodrow Wilson School - Development Studies.
- Kuznets S. 1955 Economic growth and income inequality. *Am. Econ. Rev.* **45**, 1–28.
- Harsanyi J. 1955 Cardinal welfare, individualistic ethics, and interpersonal comparisons of utility. *J. Polit. Econ.* **63**, 309–321. (doi:10.1086/257678)
- Alkire S, Foster J. 2011 Understandings and misunderstandings of multidimensional poverty measurement. *J. Econ. Inequal.* **9**, 289–314. (doi:10.1007/s10888-011-9181-4)
- Coudouel A, Hentschel J, Wodon Q. 2002 Poverty measurement and analysis.
- Babu S, Gajanan SN, Sanyal P. 2014 *Food security, poverty and nutrition policy analysis: statistical methods and applications*. New York, NY: Academic Press.
- Ravallion M. 1998 *Poverty lines in theory and practice*, pp. 1–53. Washington, DC: The World Bank.
- Caterina Ruggeri Laderchi RS, and F. S. Does it matter that we don't agree on the definition of poverty? A comparison of four approaches. Queen Elizabeth House Working Paper Series. Oxford, UK: University of Oxford.
- Falkingham J, Namazie C. 2002 *Measuring health and poverty: a review of approaches to identifying the poor*. London, UK: DFID Health Systems Resource Centre.
- Hulme D, Shepherd A. 2003 Conceptualizing chronic poverty. *World Dev.* **31**, 403–423. (doi:10.1016/S0305-750X(02)00222-X)
- Foster J, Greer J, Thorbecke E. 1984 A class of decomposable poverty measures. *Econometrica* **52**, 761–766. (doi:10.2307/1913475)
- Ligon E, Schechter L. 2003 Measuring vulnerability. *Econ. J.* **113**, C95–C102. (doi:10.1111/1468-0297.00117)
- Banerjee A, Duflo E, Chattopadhyay R, Shapiro J. 2010 Targeting the hard-core poor: an impact assessment. Poverty Action Lab. See [https://www.povertyactionlab.org/sites/default/files/publications/110-%20November%202011\\_0.pdf](https://www.povertyactionlab.org/sites/default/files/publications/110-%20November%202011_0.pdf).
- Schreiner M. 2011 Estimating expenditure-based poverty from the Bangladesh demographic and health survey — MEASURE evaluation. *Bangladesh J. Dev. Stud.* **34**, 65–94.
- Filmer D, Pritchett LH. 2001 Estimating wealth effects without expenditure data—or tears: an application to educational enrollments in states of India. *Demography* **38**, 115–132.
- Sahn DE, Stifel D. 2003 Exploring alternative measures of welfare in the absence of expenditure data. *Rev. Income Wealth* **49**, 463–489. (doi:10.1111/j.0034-6586.2003.00100.x)
- Foreit K, Schreiner M. 2011 *Comparing alternative measures of poverty: assets-based wealth index vs. expenditures-based poverty score — MEASURE evaluation*. Chapel Hill, NC: University of North Carolina at Chapel Hill.
- Liverpool-Tasie LSO, Winter-Nelson A. 2011 Asset versus consumption poverty and poverty dynamics in rural Ethiopia. *Agric. Econ.* **42**, 221–233. (doi:10.1111/j.1574-0862.2010.00479.x)
- Jerven M. 2013 *Poor numbers: how we are misled by African development statistics and what to do about it*. Ithaca, NY: Cornell University Press.
- Hentschel J, Lanjouw JO, Lanjouw P, Poggi J. 1998 *Combining census and survey data to study spatial dimensions of poverty*. Washington, DC: The World Bank Development Research Group and Poverty Reduction and Economic Management Network Poverty Division.
- Elbers C, Lanjouw JO, Lanjouw P. 2002 Micro-level estimation of poverty and inequality. *Econometrica* **71**, 355–364. (doi:10.1111/1468-0262.00399)
- Elbers C, Lanjouw JO, Lanjouw P. 2002 *Micro-level estimation of welfare*. Washington, DC: World Bank Development Research Group.
- Tatem AJ, Gething PW, Pezzullo C, Weiss D, Bhatt S. 2014 Development of high-resolution gridded poverty surfaces. See <http://www.worldpop.org.uk/resources/docs/Poverty-mapping-report.pdf>.
- Sedda L *et al.* 2015 Poverty, health and satellite-derived vegetation indices: their inter-spatial relationship in West Africa. *Int. Health* **7**, 99–106. (doi:10.1093/inthealth/ihv005)
- Pozzi F, Robinson T, Nelson A. 2009 *Accessibility mapping and rural poverty in the horn of Africa*. Rome, Italy: Food and Agriculture Organization of the United Nations.
- Robinson T, Pozzi F. 2008 *Poverty and welfare measures in the Horn of Africa*. Rome, Italy: IGAD Livestock Policy Initiative.
- Rogers D, Emwanu T, Robinson T. 2006 *Poverty mapping in Uganda: an analysis using remotely sensed and other environmental data*. Rome, Italy: Food and Agriculture Organization of the United Nations.
- Okwi PO *et al.* 2007 Spatial determinants of poverty in rural Kenya. *Proc. Natl Acad. Sci. USA* **104**, 16 769–16 774. (doi:10.1073/pnas.0611107104)
- Jean N *et al.* 2016 Combining satellite imagery and machine learning to predict poverty. *Science* **353**, 790–794. (doi:10.1126/science.aaf7894)
- Blumenstock J, Cadamuro G. 2015 On R. Predicting poverty and wealth from mobile phone metadata. *Science* **350**, 1073–1076.
- Smith-Clarke C, Mashhadi A, Capra L. 2014 Poverty on the cheap: estimating poverty maps using aggregated mobile communication networks. In *Proc. of the SIGCHI Conf. on Human Factors in Computing Systems, Toronto, Ontario, Canada*, pp. 511–520. New York, NY: ACM.
- Frias-Martinez V, Virseda J. 2012 On the relationship between socio-economic factors and cell phone usage. In *Proc. of the Fifth Int. Conf. on Information and Communication Technologies and Development*, Atlanta, GA, pp. 76–84. New York, NY: ACM.
- Soto V, Frias-Martinez V, Virseda J, Frias-Martinez E. 2011 Prediction of socioeconomic levels using cell phone records. In *User modeling, adaption and personalization* (eds JA Konstan, R Conejo, JL Marzo, N Oliver), pp. 377–388. Berlin, Heidelberg: Springer.
- Eagle N, Macy M, Claxton R. 2010 Network diversity and economic development. *Science* **328**, 1029–1031. (doi:10.1126/science.1186605)
- Blumenstock JE, Eagle N. 2012 Divided we call: disparities in access and use of mobile phones in Rwanda. *Inf. Technol. Int. Dev.* **8**, 1–16.
- Desiere S, Vellema W, D'Haese M. 2015 A validity assessment of the Progress out of Poverty Index (PPI)<sup>TM</sup>. *Eval. Program Plann.* **49**, 10–18. (doi:10.1016/j.evalprogplan.2014.11.002)
- Okabe A, Boots B, Sugihara K, Chiu SN. 2009 *Spatial tessellations: concepts and applications of voronoi diagrams*. Chichester, UK: John Wiley & Sons.
- Intermedia. 2014 Intermedia Financial Inclusion Insight Project Bangladesh - Steps Toward Financial Inclusion 2014. See <http://finclusion.org/>.
- Rutstein S. 2008 *The DHS wealth index: approaches for rural and urban areas*. Calverton, MD: United States Agency for International Development.
- Schreiner M. 2013 A Simple Poverty Scorecard for Bangladesh. See [http://microfinance.com/English/Papers/Scoring\\_Poverty\\_Bangladesh\\_2010\\_EN.pdf](http://microfinance.com/English/Papers/Scoring_Poverty_Bangladesh_2010_EN.pdf).
- Grameenphone, Bangladesh. 2015 Telenor Group. See <http://www.telenor.com/investors/company-facts/business-description/grameenphone-bangladesh/> (accessed: 15 November 2015).
- Noor AM, Alegana VA, Gething PW, Tatem AJ, Snow RW. 2008 Using remotely sensed night-time light as a proxy for poverty in Africa. *Popul. Health Metr.* **6**, 5. (doi:10.1186/1478-7954-6-5)
- Ghosh T, Anderson SJ, Elvidge CD, Sutton PC. 2013 Using nighttime satellite imagery as a proxy

- measure of human well-being. *Sustainability* **5**, 4988–5019. (doi:10.3390/su5124988)
50. Watmough GR, Atkinson PM, Saikia A, Hutton CW. 2016 Understanding the evidence base for poverty–environment relationships using remotely sensed satellite data: an example from Assam, India. *World Dev.* **78**, 188–203. (doi:10.1016/j.worlddev.2015.10.031)
  51. Murtaugh PA. 2009 Performance of several variable-selection methods applied to real ecological data. *Ecol. Lett.* **12**, 1061–1068. (doi:10.1111/j.1461-0248.2009.01361.x)
  52. glmulti: An R package for easy automated model selection with (generalized) linear models. *J. Stat. softw.* See <http://www.jstatsoft.org/article/view/v034i12> (accessed: 21 January 2016).
  53. R Core Team. 2015 *R: a language and environment for statistical computing*. Vienna, Austria: R Foundation for Statistical Computing.
  54. Blangiardo M, Cameletti M, Baio G, Rue H. 2013 Spatial and spatio-temporal models with R-INLA. *Spat. Spatio-Temporal Epidemiol.* **4**, 33–49. (doi:10.1016/j.sste.2012.12.001)
  55. Blangiardo M, Cameletti M. 2015 *Spatial and spatio-temporal Bayesian models with R - INLA*. New York, NY: John Wiley & Sons.
  56. Sørbye SH, Rue H. 2014 Scaling intrinsic Gaussian Markov random field priors in spatial modelling. *Spat. Stat.* **8**, 39–51. (doi:10.1016/j.spasta.2013.06.004)
  57. Besag J, Kooperberg C. 1995 On conditional and intrinsic autoregressions. *Biometrika* **82**, 733–746.
  58. Rue H, Martino S, Chopin N. 2009 Approximate Bayesian inference for latent Gaussian models by using integrated nested Laplace approximations. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **71**, 319–392. (doi:10.1111/j.1467-9868.2008.00700.x)
  59. Rue H, Martino S. 2007 Approximate Bayesian inference for hierarchical Gaussian Markov random field models. *J. Stat. Plan. Inference* **137**, 3177–3192. (doi:10.1016/j.jspi.2006.07.016)
  60. The R-INLA project. Latent models. The R-INLA project. 2016 See <http://www.r-inla.org/models/latent-models> (accessed: 21st January 2016).
  61. The R-INLA project. 2015 Besag model for spatial effects. See <http://www.math.ntnu.no/inla/r-inla.org/doc/latent/besag.pdf>.
  62. Spiegelhalter DJ, Best NG, Carlin BP, Van Der Linde A. 2002 Bayesian measures of model complexity and fit. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **64**, 583–639. (doi:10.1111/1467-9868.00353)
  63. Gruebner O *et al.* 2014 Mapping the slums of Dhaka from 2006 to 2010. *Dataset Pap. Sci. Dataset Pap. Sci.* **2014**, e172182.
  64. Ahmad N *et al.* 2010 *Technical report*, pp. 1–56. Singapore: The World Bank.
  65. Blondel VD, Decuyper A, Krings G. 2015 A survey of results on mobile phone datasets analysis. *EPJ Data Sci.* **4**, 10. (doi:10.1140/epjds/s13688-015-0046-0)
  66. Barrett CB. 2005 Rural poverty dynamics: development policy implications. *Agric. Econ.* **32**, 45–60. (doi:10.1111/j.0169-5150.2004.00013.x)
  67. Bedi T, Coudouel A, Simler K. 2007 *More than a pretty picture: using poverty maps to design better policies and interventions*. Washington, DC: World Bank Publications.
  68. Toole JL *et al.* 2015 Tracking employment shocks using mobile phone data. *J. R. Soc. Interface* **12**, 20150185. (doi:10.1098/rsif.2015.0185)



# Mapping poverty using mobile phone and satellite data

J.E. Steele, P. Sundsoy, C. Pezzulo, V. Alegana, T. Bird, J. Blumenstock, J. Bjelland, K. Engo-Monsen, YA de Montjoye, A. Iqbal, K. Hadiuzzaman, X. Lu, E. Wetter, A.J. Tatem, and L. Bengtsson

## Supplementary Information (SI)

### Table of Contents

A. Input data .....	3
A.1. Geolocated survey data .....	3
A.1.1. Demographic and Health Survey Wealth Index .....	3
A.1.2. Progress out of Poverty Index .....	4
A.1.3. Market research household surveys and income data .....	4
A.2. Mobile phone call detail records (CDRs) .....	5
A.3. Remote Sensing-GIS covariates .....	6
A.3.1 GPS data displacement .....	6
B. Statistical analyses and prediction mapping .....	7
B.1. Covariate selection via generalized linear models .....	7
B.2. Prediction mapping via Bayesian geostatistical models .....	7
References .....	8
SI Tables and Figures .....	11
Table S1. Summary information for remote sensing-GIS and mobile phone call detail record datasets used for covariate selection and Bayesian geostatistical poverty mapping. ....	11
Table S2A. Wealth Index models for RS-only, CDR-only, and CDR+RS data: national, urban, and rural strata. ....	13
Table S2B. Progress out of Poverty Index models for RS-only, CDR-only, and CDR+RS data: national, urban, and rural strata. ....	14
Table S2C. Income models for RS-only, CDR-only, and CDR+RS data: national, urban, and rural strata. ....	15
Table S3. Comparison of $r^2$ and RMSE for INLA models run with only a structured spatial random effect (Spatial interpolation) and the full model (Spatial model + covariates). ....	16
Table S4. Comparison of deviance information criterion (DIC) model fit for Bayesian geostatistical models run with a structured spatial random effect (Spatial model) and without (Non-spatial model). ....	17
Figure S1. Survey sample locations for DHS wealth index (A), Progress out of Poverty Index (B), and income survey respondents (C). ....	18
Figure S2. National level prediction maps for mean wealth index (A) with uncertainty (D); mean probability of households being below \$2.50/day (B) with uncertainty (E); and mean USD income (C) with uncertainty (F). Maps were generated using call detail record features only and Bayesian geostatistical models. Red indicates poorer areas in prediction maps, and higher error in uncertainty maps. ....	19
Figure S3. National level prediction maps for mean wealth index (A) with uncertainty (D); mean probability of households being below \$2.50/day (B) with uncertainty (E); and mean USD income (C) with uncertainty (F). Wealth index and income maps were generated using remote sensing data only; PPI maps were generated using call detail record features and remote sensing data. All maps were generated using Bayesian geostatistical models. Red indicates poorer areas in prediction maps, and higher error in uncertainty maps. ....	20

Figure S4. Stratified urban/rural prediction maps for mean wealth index (A) with uncertainty (D); mean probability of households being below \$2.50/day (B) with uncertainty (E); and mean USD income (C) with uncertainty (F). Maps were generated using call detail record features only and Bayesian geostatistical models. Red indicates poorer areas in prediction maps, and higher error in uncertainty maps. Figure S5. Stratified urban/rural prediction maps for mean wealth index (A) with uncertainty (D); mean probability of households being below \$2.50/day (B) with uncertainty (E); and mean USD income (C) with uncertainty (F). Maps were generated using call detail record features, remote sensing data, and Bayesian geostatistical models. Red indicates poorer areas in prediction maps, and higher error in uncertainty maps. ....21

Figure S5. Stratified urban/rural prediction maps for mean wealth index (A) with uncertainty (D); mean probability of households being below \$2.50/day (B) with uncertainty (E); and mean USD income (C) with uncertainty (F). Maps were generated using call detail record features, remote sensing data, and Bayesian geostatistical models. Red indicates poorer areas in prediction maps, and higher error in uncertainty maps. ....22

Figure S6. Stratified urban/rural prediction maps for mean wealth index (A) with uncertainty (D); mean probability of households being below \$2.50/day (B) with uncertainty (E); and mean USD income (C) with uncertainty (F). Maps were generated using remote sensing data only and Bayesian geostatistical models. Red indicates poorer areas in prediction maps, and higher error in uncertainty maps. ....23

## A. Input data

### A.1. Geolocated survey data

The growing number of georeferenced household survey data from low and middle-income countries allows us to explore poverty metrics and comparisons between them while explicitly considering their geographic distribution. In Bangladesh, we utilised three geographically referenced datasets representing asset, consumption, and income-based measures of wellbeing (Figure S1).

#### A.1.1. Demographic and Health Survey Wealth Index

The Demographic and Health Surveys (DHS) were designed primarily to collect household data on marriage, fertility, family planning, reproductive and child health, and HIV/AIDS in almost all lower income countries<sup>1</sup>. Through the assembling of indicators correlated with a household's economic status (e.g. ownership of television, telephone, radio as well as variables describing type of floor and ceiling material and other facilities), a wealth index is calculated for each country at each time<sup>2</sup> based on the idea that the possession of assets and access to services and amenities are related to the relative economic position of the household in the country<sup>3</sup>. By its construction, the wealth index is a relative measure of wealth within each survey; however, a new methodology has been developed in order to make it comparable across countries and through time<sup>4</sup>. Moreover, recent adjustments have been made to the methods of constructing the wealth index to overcome criticism that the original score was not adequately capturing the differences between urban and rural poverty or identifying the poorest of the poor<sup>3</sup>.

The wealth index is constructed using a principal component analysis (PCA), which includes a long list of assets owned by households as well as other indicators. (The complete list of indicators included in the PCAs for each survey, as well as PCA analysis and results can be found at <http://dhsprogram.com/topics/wealth-index/Wealth-Index-Construction.cfm>). The first factor from the PCA, capturing the largest percentage of the variance within the dataset, is derived adjusting for urban and rural strata<sup>3,5</sup>. In practice, a national index and two area-specific indexes representing urban and rural strata are individually constructed using sets of assets/services specific to each in order to better capture differences between urban and rural areas, and compare the wealth index between them<sup>3</sup>. Subsequently, applying regression techniques described in Rutstein<sup>3</sup> and Rutstein<sup>6</sup>, the three indexes are combined into a single wealth distribution and a composite national index is derived. This method ensures comparability between urban and rural areas.

Here we used the 2011 Bangladesh DHS<sup>7</sup> (Figure S1A), a nationally representative survey based on a two-stage stratified sample of households, where 600 enumeration areas (EA or cluster) were first selected with probability proportional to the EA size, (207 clusters in urban areas and 393 in rural areas). This first stage of selection provided a listing of households for the second stage, where a systematic sample of 30 households on average was selected per cluster, to create statistically reliable estimates of key demographic and health variables<sup>7,8</sup>. In recent DHS surveys where HIV/AIDS data are not collected, geolocations for each cluster are available. The survey cluster coordinates represent an estimated centre of the cluster and are collected in the field through GPS receivers. To maintain respondents' confidentiality, GPS positions for all clusters are randomly displaced by a maximum of five kilometres for rural clusters and a maximum of two kilometres for urban clusters<sup>9–11</sup>.

#### A.1.2. Progress out of Poverty Index

The Progress out of Poverty Index (PPI) (Figure S1B) was designed to be easily collected, simple and cost-effective to implement and verify<sup>12</sup>, while applying a rigorous methodology through selecting assets based on their statistical relationship with poverty<sup>12,13</sup>. In the case of Bangladesh, an easy-to-use poverty scorecard<sup>13</sup> of 10 questions was created in March 2013, based on data from the Bangladesh 2010 Household Income and Expenditure Survey (HIES). The questions selected are aggregated into a score highly correlated with poverty status as measured by the HIES. The scores included in the scorecard are then translated into likelihoods that the household has per-capita expenditure above or below a given poverty line<sup>13</sup>.

A nationally representative survey of all adults in Bangladesh was undertaken by InterMedia Financial Inclusion Insight Project ([www.finclusion.org](http://www.finclusion.org)) in 2014 (wave 2), where 6,000 Bangladeshi individuals aged 15 and above were interviewed<sup>14</sup> and geolocations for each individual were included in the survey. In Bangladesh, InterMedia adopted a stratified sample strategy, whereby divisions and subdivisions were first identified and interviews within each subdivision were distributed in proportion to population size. In order to select the individuals to interview, households were first randomly selected using electoral rolls to randomly assign starting points in each selected subdivision. After having identified the starting point, subsequent households were selected using the right-hand rule, and the Kish Grid method was applied to select an individual respondent from each household<sup>14,15</sup>.

#### A.1.3. Market research household surveys and income data

Two sequential large-scale market research household surveys were run by Telenor through its subsidiary, Grameenphone (GP), during 2 time periods between November and December 2013 (N=82,834, of which 55.3% GP subscribers) and February and March 2014 (N=87,509, of which 54.5% GP subscribers) (Figure S1C). The country was stratified in 226 sales territories by the phone company, and for every territory, an equal number of unions (in rural areas) and wards (in urban areas) were randomly selected. Four hundred households were surveyed in each territory, where a household was defined as a group of people sharing food from the same chula (fire/gas burner) or living under the same roof. Systematic sampling was then undertaken to select households by selecting every fourth household, starting from the selection of a random geographic point and direction within each ward or union. In the case of more than one household present in the complex or building, the fourth household was selected. In cases of non-response, the next household was then selected. Non-response rate was approximately 10% of households. Respondents within the household were selected via the Kish grid method<sup>15</sup> among those who were eligible. Eligibility was defined as individuals with their own phone, between 15 and 65 years of age. If a phone was shared between family members, usually the male head of household was interviewed. When the selected person was not home, the surveyors returned multiple times to try to reach the selected person. A very low non-response rate (less than 1-2%) was detected among respondents. The surveys were undertaken during working hours. To avoid that too many housewives were interviewed, given that men are more likely away for work, a ceiling on the number of housewives who could participate was also established. Sampling weights were applied to ensure national representativeness and correct for population sizes in urban and rural areas. Data quality control mechanisms were implemented and undertaken by the company; however, some sources of error were detected in matching household locations to phone number (approximately 20% of the cases).



## A.2. Mobile phone call detail records (CDRs)

For the household income survey respondents described above, we collected 3-months of mobile phone metadata by subscriber consent. These metadata included call detail records (CDR) and top-up information, which were further processed into features. For each survey participant, 150 features from seven different feature families were constructed (Table S1). Household income was then linked to these metadata, resulting in three months of phone usage, matched with household income for each survey respondent. To preserve user anonymity, the local operator removes all personally identifying information from the data before analysis.

To be able to map poverty in other countries we focused on features that are easily reproducible, and easy to implement by local data warehouses. Most mobile operators generate similar features. CDR features range from metrics such as basic phone usage, top-up pattern, and social network to metrics of user mobility and handset usage. They include various parameters of the corresponding distributions such as weekly or monthly median, mean, and variance. In addition, we received pre-aggregated datasets of tower-level activity from 48,190,926 subscriber SIMs over a 4-month period. This includes monthly number of subscribers per home cell, where home cell corresponds to most frequent tower. These per-user features are not directly used, but further aggregated to the Voronoi polygons, and the aggregate features are used in covariate selection, model fitting, and prediction.

At the time of data acquisition, the mobile phone operator had an approximate 42% market share, and was the largest provider of mobile telecommunication services in Bangladesh. Multi-SIM activity is common in Bangladesh, but we believe that this should not create a systematic bias in poverty estimates because the geographic coverage of the operator is so extensive. In order to comply with national laws and regulations of Bangladesh, and the privacy policy of the Telenor group, the following measures were implemented in order to preserve the privacy rights of Grameenphone customers:

- 1) All customers are de-identified and only Telenor/Grameenphone employees have had access to any detailed CDR-/top-up data;
- 2) The processing of detailed CDR/top-up data resulted in aggregations of the data on a tower-level granularity; the tower-level aggregation makes re-identification impossible.

Hence, the resulting aggregated dataset is truly anonymous and involves no personal data.

Compared with other countries of comparable income levels, Bangladesh has a high mobile phone penetration, which includes rural areas. Fifty percent of the population above the age of fifteen has a mobile subscription<sup>16</sup>. The proportion of households with at least one mobile phone is increasing rapidly; between 2011 and 2014, household ownership across the whole of Bangladesh rose from 78% to 89%, with much of that growth concentrated among rural households<sup>17</sup>. The CDR data used in this study are available upon request for the replication of results only by contacting the corresponding author.

### A.3. Remote Sensing-GIS covariates

Ancillary data layers used as remote sensing-GIS (hereafter RS) covariates were identified, assembled, and processed for the whole of Bangladesh at a 1-km spatial resolution. These data are described in Table S1 and include 25 raster and vector datasets obtained from existing sources or produced ad hoc for this study to include environmental and physical metrics likely to be associated with human welfare<sup>18–22</sup>. These data differed in spatial and temporal resolution, type, accuracy, and coverage. In order to align all data for model fitting and prediction, the following steps were taken:

- 1) Bangladesh was rasterized at a resolution of 30-arcsec (0.00833333 degree, corresponding to approximately 1-km at the equator);
- 2) Vector datasets were rasterized at a resolution of 30-arcsec;
- 3) When necessary, raster datasets were resampled to a resolution of 30-arcsec using an interpolation technique appropriate for the resolution and type of the original dataset;
- 4) All datasets were spatially aligned to make every pixel representing the same location coincident and match the rasterized study area.

Furthermore, for ad hoc datasets such as distance to roads and waterways, we used a customized Azimuthal Equidistant projection centred in the middle of the study area and clipped to a buffer extending 100 metres beyond its boundary to project the input data. This buffered area was rasterized to a resolution of approximately 927 metres, corresponding to 30-arcsec at the centre of the study area where distortion is smallest. Euclidean distance was calculated for each distance-to-covariate within the customized projection. The resultant layers were then projected back to GCS WGS84, and made coincident with the rasterized study area. All datasets representing categorical variables (e.g. protected areas, global urban extent, etc.) were projected, rasterized, and/or resampled to 1-km resolution, spatially aligned to the rasterized study area, and converted into binary covariates, representing the presence or absence of a given feature. This resulted in twenty-five 1-km raster datasets, which were used to extract the mean, mode, or sum of each covariate for each Voronoi polygon, dependent on the type of dataset. These values were used for covariate selection, model fitting, and prediction.

#### A.3.1 GPS data displacement

In addition to the aforementioned processing, additional steps were undertaken to appropriately account for the displacement inherent in DHS data. When these data are collected, the latitude and longitude of the centre of each DHS cluster (representing numerous households) is collected in the field with a GPS receiver. To maintain respondents' confidentiality, GPS latitude/longitude positions for all DHS clusters are randomly displaced by a maximum of five kilometres for rural clusters and two kilometres for urban clusters. The displacement is restricted so the points stay within the country, within the DHS survey region, and within the second administrative level<sup>9–11</sup>.

In order to account for the displacement in our analyses, we created buffers around each cluster centroid of 2 km and 5 km for urban and rural clusters, respectively, and subsequently extracted the RS covariate data for each buffer zone. For continuous covariates, the minimum, maximum, and mean values were calculated and extracted. For categorical covariates, the modal value was calculated and extracted.

## B. Statistical analyses and prediction mapping

### B.1. Covariate selection via generalized linear models

Stratifying models into urban and rural components produced the best fit models as measured by AIC. Top-up data produced the most important CDR feature family for all poverty measures and models. Within this feature family, significant covariates included recharge amounts and frequencies per tower, spending speeds and time between refills, and fractions of the lowest and highest available top-up amounts. Advanced phone usage was also an important CDR feature family, especially for PPI and income models. Sum, revenue, count, and volume of incoming and outgoing multimedia messaging, Internet usage, and videos were prominent. Basic phone usage covariates measuring incoming and outgoing text counts were important for every model save for rural WI models. Mobility covariates including number and entropy of places and radius of gyration were also significant features in all three strata and poverty measures, as were social network features such as number and entropy of contacts.

Nighttime lights and covariates representing access - especially transport time to closest urban settlement and distance to roads - were the most important RS covariates for all three poverty measures and strata. Vegetation productivity, as measured by the Enhanced Vegetation Index (EVI), and elevation were also prominent RS features in all three strata, whereas climate variables featured prominently in rural models.

### B.2. Prediction mapping via Bayesian geostatistical models

Using the models selected and described in Tables S2A-C, we employed hierarchical Bayesian geostatistical models (BGMs) for prediction as described in our manuscript. All prediction maps not highlighted in our manuscript can be found in Figures S2-S6. Model performance was based on out-of-sample validation statistics calculated on a 20% test subset of poverty data input points (see Table 1 in manuscript). The performance of models built with CDR-only or RS-only data varied based on poverty measure and strata. RS-only models were more successful at predicting the WI for all three strata ( $r^2 = 0.74, 0.71, 0.72$  for national, urban, and rural models), as compared to CDR-only models. However, the CDR-only models performed nearly as well ( $r^2 = 0.64, 0.70, 0.50$  for national, urban, and rural models), and all urban WI models including CDRs outperformed national level models. The urban CDR-RS model exhibits the highest explained variance for any model ( $r^2 = 0.78$ ), and the urban CDR-only model outperforms the national CDR-only model ( $r^2 = 0.70$  versus  $r^2 = 0.64$ , respectively). For PPI and income measures of poverty, CDR data produced the best models in urban areas, whereas RS data produced the best models in rural areas. This highlights the compatibility of these two datasets for predicting different measures of poverty at different scales, as the best estimates and lowest error corresponded to the data with fine-scale spatial heterogeneity (CDRs in urban areas; RS data in rural areas). To that end, national poverty models generally performed best when utilising both CDRs and RS data.

To compare full model performance against a spatial interpolation model, we modelled the training data for all three poverty indicators using only the spatial random effect in the INLA model (see section 2.5 in manuscript). These results are shown in Table S3. We compared out-of-sample  $r^2$  and RMSE values against results from the full models (see Table 1 in manuscript). The results show a spatial pattern in the WI data as the model built with only a spatial random effect yields an  $r^2 = 0.49$ , RMSE = 0.578. When compared to the full model, the addition of covariate data increases the  $r^2$  to 0.76, and the RMSE decreases to 0.394. Similarly for income, the data do show a slight spatial pattern, but the addition of covariate data to the model increases the predictive power and

decreases the error. For the PPI, the covariates do not show a strong influence in the modelling results, and the model was driven by the spatial process, which suggests there's an underlying spatial covariate that we're not capturing in the model that could explain the data.

Model fit based on the spatial effect can also be considered using DIC, a hierarchical modelling generalization of the AIC and BIC, which can be useful in Bayesian modelling comparison. The BIC allows for comparing models using criterion based on the trade-off between the fit of the data to the model and the corresponding complexity of the model. Models with smaller DIC values are preferred over models with larger DIC values as the measure favours better fit and fewer parameters<sup>23</sup>. These results are shown in Table S4. For nearly every model with CDR data, DIC is greatly improved by accounting for the spatial covariance in the data structure. However, the income models see slight or no improvement from including the random spatial effect, likely due to the fact that they include and are thus penalised for many covariates.

## References

1. Rutstein, S. & Rojas, G. *Guide to DHS Statistics*. 168 (ORC Macro, Demographic and Health Surveys, 2006).
2. Rutstein, S. & Johnson, K. *The DHS Wealth Index*. (ORC Macro, Demographic and Health Surveys, 2004).
3. Rutstein, S. *The DHS Wealth Index: Approaches for Rural and Urban Areas*. (United States Agency for International Development, 2008).
4. Rutstein, S. & Staveteig, S. *Making the Demographic and Health Surveys Wealth Index Comparable*. (ICF International).
5. Rogers, D., Wint, G., Alexander, N., Pozzi, F. & Robinson, T. *Wealth Index Mapping in the Horn of Africa*. (Food and Agriculture Organization of the United Nations, 2011).
6. Rutstein, S. *Steps to constructing the new DHS Wealth Index*.
7. National Institute of Population Research and Training, Mitra and Associates & ICF International. *Bangladesh Demographic and Health Survey 2011*. (2013).
8. ICF International. *Demographic and Health Survey Sampling Household Listing Manual*. (ICF International, 2012).

9. Burgert, C., Colston, J., Thea Roy, T. & Zachary, B. *Geographic displacement procedure and georeferenced data release policy for the Demographic and Health Surveys*. (ICF International, 2013).
10. Perez-Heydrick, C., Warren, J., Burgert, C. & Emch, M. *Guidelines on the Use of DHS GPS Data*. (ICF International, 2013).
11. DHS Spatial Interpolation Working Group. *Spatial Interpolation with Demographic and Health Survey Data: Key Considerations*. (ICF International, 2014).
12. Desiere, S., Vellema, W. & D'Haese, M. A validity assessment of the Progress out of Poverty Index (PPI)<sup>TM</sup>. *Eval. Program Plann.* **49**, 10–18 (2015).
13. Schreiner, M. *A Simple Poverty Scorecard for Bangladesh*. (2013).
14. Intermedia. *Intermedia Financial Inclusion Insight Project Bangladesh - Steps Toward Financial Inclusion 2014*. (2014).
15. Kish, L. A Procedure for Objective Respondent Selection within the Household. *J. Am. Stat. Assoc.* **44**, 380–387 (1949).
16. Lucini, B. A. & Hatt, T. *Country overview: Bangladesh*. (GSMA Intelligence, 2014).
17. National Institute of Population Research and Training, Mitra and Associates & The DHS Program. *Bangladesh Demographic and Health Survey 2014: Key Indicators*. (NIPORT, Mitra and Associates, and ICF International, 2015).
18. Pozzi, F., Robinson, T. & Nelson, A. *Accessibility Mapping and Rural Poverty in the Horn of Africa*. (Food and Agriculture Organization of the United Nations, 2009).
19. Rogers, D., Emwanu, T. & Robinson, T. *Poverty Mapping in Uganda: An Analysis Using Remotely Sensed and Other Environmental Data*. (Food and Agriculture Organization of the United Nations, 2006).
20. Noor, A. M., Alegana, V. A., Gething, P. W., Tatem, A. J. & Snow, R. W. Using remotely sensed night-time light as a proxy for poverty in Africa. *Popul. Health Metr.* **6**, 5 (2008).

Mapping poverty using mobile phone and satellite data

*J.E. Steele, P. Sundsoy, C. Pezzulo, V. Alegana, T. Bird, J. Blumenstock, J. Bjelland, K. Engo-Monsen, YA de Montjoye, A. Iqbal, K. Hadiuzzaman, X. Lu, E. Wetter, A.J. Tatem, and L. Bengtsson*

21. Ghosh, T., Anderson, S. J., Elvidge, C. D. & Sutton, P. C. Using Nighttime Satellite Imagery as a Proxy Measure of Human Well-Being. *Sustainability* **5**, 4988–5019 (2013).
22. Watmough, G. R., Atkinson, P. M., Saikia, A. & Hutton, C. W. Understanding the Evidence Base for Poverty–Environment Relationships using Remotely Sensed Satellite Data: An Example from Assam, India. *World Dev.* **78**, 188–203 (2016).
23. Spiegelhalter, D. J., Best, N. G., Carlin, B. P. & Van Der Linde, A. Bayesian measures of model complexity and fit. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **64**, 583–639 (2002).



## SI Tables and Figures

Table S1. Summary information for remote sensing-GIS and mobile phone call detail record datasets used for covariate selection and Bayesian geostatistical poverty mapping.

Category	Description	Source	Resolution (Degrees)	Year
<b>RS-GIS</b>				
Accessibility	Accessibility to populated places with more than 50k people	European Commission Joint Research Centre ( <a href="http://forobs.jrc.ec.europa.eu/products/gam/">http://forobs.jrc.ec.europa.eu/products/gam/</a> )	0.0083333	2000
Population	Population count [per pixel]	WorldPop ( <a href="http://www.worldpop.org.uk/">http://www.worldpop.org.uk/</a> )	0.0008333	2010
Population	Population count [per pixel]	CIESIN - Global Rural Urban Mapping Project ( <a href="http://sedac.ciesin.columbia.edu/data/collection/grump-v1/sets/browse">http://sedac.ciesin.columbia.edu/data/collection/grump-v1/sets/browse</a> )	0.0083333	2000
Population	Population density [per sqkm]	CIESIN - Global Rural Urban Mapping Project ( <a href="http://sedac.ciesin.columbia.edu/data/collection/grump-v1/sets/browse">http://sedac.ciesin.columbia.edu/data/collection/grump-v1/sets/browse</a> )	0.0083333	2000
Climate	Mean Aridity Index	CGIAR-CSI ( <a href="http://www.cgiar-csi.org/data">http://www.cgiar-csi.org/data</a> )	0.0083333	1950-2000
Climate	Average annual Potential Evapotranspiration [mm]	CGIAR-CSI ( <a href="http://www.cgiar-csi.org/data">http://www.cgiar-csi.org/data</a> )	0.0083333	1950-2000
Night-time lights	VIIRS night-time lights [W cm-2 sr-1]	NOAA VIIRS ( <a href="http://ngdc.noaa.gov/eog/viirs.html">http://ngdc.noaa.gov/eog/viirs.html</a> )	0.0041667	2014
Elevation	Elevation [meter]	CGIAR-CSI ( <a href="http://srtm.cgiar.org/">http://srtm.cgiar.org/</a> )	0.0083316	2008
Vegetation	Vegetation	MODIS MOD13A1 [Enhanced vegetation index]	0.0041667	2010-2014
Distance	Distance to roads [meter]	Input data from OSM ( <a href="http://extract.bbbike.org/">http://extract.bbbike.org/</a> )	0.0083333	2014
Distance	Distance to waterways [meter]	Input data from OSM ( <a href="http://extract.bbbike.org/">http://extract.bbbike.org/</a> )	0.0083333	2014
Urban/Rural	Urban/Rural	MODIS-based Global Urban extent	0.0041670	2000-2001
Urban/Rural	Urban/Rural	CIESIN - Global Rural Urban Mapping Project ( <a href="http://sedac.ciesin.columbia.edu/data/collection/grump-v1/sets/browse">http://sedac.ciesin.columbia.edu/data/collection/grump-v1/sets/browse</a> )	0.0083333	2000
Urban/Rural	Global Human Settlement Layer	Global Land Cover Facility ( <a href="http://www.landcover.org">www.landcover.org</a> )	0.002818	2014
Protected Area	Protected areas	WDPA ( <a href="http://www.protectedplanet.net/">http://www.protectedplanet.net/</a> )	Vector	2012
Land Cover	Land cover	ESA GlobCover Project ( <a href="http://due.esrin.esa.int/page_globcover.php">http://due.esrin.esa.int/page_globcover.php</a> )	0.0027777	2009
Land Cover	Land cover	IGBP MODIS MCD12Q1 ( <a href="https://lpdaac.usgs.gov/dataset_discovery/modis/modis_products_table/mcd12q1">https://lpdaac.usgs.gov/dataset_discovery/modis/modis_products_table/mcd12q1</a> )	0.0041670	2012
Land Cover	Land cover	ONRL DAAC Synergetic Land Cover Product (SYNMAP) ( <a href="http://webmap.ornl.gov/wcsdown/wcsdown.jsp?dg_id=10024_1">http://webmap.ornl.gov/wcsdown/wcsdown.jsp?dg_id=10024_1</a> )	0.0083333	2000-2001

# Mapping poverty using mobile phone and satellite data

J.E. Steele, P. Sundsoy, C. Pezzulo, V. Alegana, T. Bird, J. Blumenstock, J. Bjelland, K. Engo-Monsen, YA de Monjoye, A. Iqbal, K. Hadiuzzaman, X. Lu, E. Wetter, A.J. Tatem, and L. Bengtsson

Demographic	Pregnancies		
Demographic	Births	WorldPop ( <a href="http://www.worldpop.org.uk/">http://www.worldpop.org.uk/</a> )	0.0008333 2012
Ethnicity	Georeferenced ethnic groups	WorldPop ( <a href="http://www.worldpop.org.uk/">http://www.worldpop.org.uk/</a> )	0.0008333 2012
Climate	Mean annual precipitation	ETH Zurich ( <a href="http://www.icr.ethz.ch/data/geoepr">http://www.icr.ethz.ch/data/geoepr</a> )	Vector 2014
Climate	Mean annual temperature	WorldClim ( <a href="http://www.worldclim.org/download">http://www.worldclim.org/download</a> )	0.0083333 1950-2000
		WorldClim ( <a href="http://www.worldclim.org/download">http://www.worldclim.org/download</a> )	0.0083333 1950-2000

## Call Detail Records

Basic phone usage	Outgoing/incoming voice duration, SMS count, etc.	Telenor/Grameenphone	NA	2013-2014
Top-up transactions	Spending speed, recharge amount per transaction, fraction of lowest/highest recharge amount, coefficient of variation recharge amount, etc.	Telenor/Grameenphone	NA	2013-2014
Location/mobility	Home district/tower, radius of gyration, entropy of places, number of places, etc.	Telenor/Grameenphone	NA	2013-2014
Social Network	Interaction per contact, degree, entropy of contacts, etc.	Telenor/Grameenphone	NA	2013-2014
Handset type	Brand, manufacturer, camera enabled, smart/feature/basic phone, etc.	Telenor/Grameenphone	NA	2013-2014
Revenue	Charge of outgoing/incoming SMS, MMS, voice, video, value added services, roaming, internet, etc.	Telenor/Grameenphone	NA	2013-2014
Advanced phone usage	Internet volume/count, MMS count, video count/duration, value added services duration/count, etc.	Telenor/Grameenphone	NA	2013-2014



# Mapping poverty using mobile phone and satellite data

J.E. Steele, P. Sundsoy, C. Pezzulo, V. Alegana, T. Bird, J. Blumenstock, J. Bjelland, K. Engo-Monsen, YA de Monjoye, A. Iqbal, K. Hadiuzzaman, X. Lu, E. Wetter, A.J. Tatem, and L. Bengtsson

Table S2A. Wealth Index models for RS-only, CDR-only, and CDR+RS data: national, urban, and rural strata.

AIC	NATIONAL	URBAN	RURAL
RS-only	690.61	333.44	161.42
CDR-only	907.78	373.97	164.14
CDR-RS	651.76	318.12	115.52
<b>MODEL</b>			
RS-only	1 + transport time to closest urban settlement + nighttime lights + EVI + elevation	1 + distance to roads + distance to waterways + nighttime lights + elevation	1 + transport time to closest urban settlement + annual temperature + annual precipitation + distance to roads + distance to waterways + nighttime lights
CDR-only	1 + recharge average per tower + percent nocturnal calls + number of places + entropy of contacts + outgoing internet sessions + sum outgoing internet sessions + incoming voice duration + count incoming content management system + count sum incoming content management system + volume of incoming multimedia messages + recharge amount per transaction + count incoming multimedia messages + count incoming texts + weekly recharge amount	1 + recharge average per tower + number of places + entropy of contacts + spending speed + average outgoing text count + sum count incoming content management system + weekly recharge amount	1 + recharge average per tower + percent nocturnal calls + entropy of places + radius of gyration + interactions per contact + recharge amount (CV) + number of retailers visited weekly (CV) + sum incoming video duration + count incoming multimedia messages + weekly recharge frequency (CV) + sum incoming video count + recharge amount per transaction (CV)
CDR-RS	1 + transport time to closest urban settlement + nighttime lights + EVI + elevation + recharge average per tower + percent nocturnal calls + outgoing internet sessions + count incoming content management system + recharge amount per transaction + count incoming texts + weekly recharge amount	1 + distance to roads + distance to waterways + nighttime lights + elevation + recharge average per tower + spending speed + average outgoing text count + weekly recharge amount	1 + transport time to closest urban settlement + annual temperature + distance to roads + distance to waterways + nighttime lights + recharge average per tower + percent nocturnal calls + entropy of places + radius of gyration + interactions per contact + recharge amount (CV) + number of retailers visited weekly (CV) + sum incoming video duration + weekly recharge frequency (CV) + sum incoming video count + recharge amount per transaction (CV)

\*CV=coefficient of variation

# Mapping poverty using mobile phone and satellite data

*J.E. Steele, P. Sundsoy, C. Pezzulo, V. Alegana, T. Bird, J. Blumenstock, J. Bjelland, K. Engo-Monsen, YA de Monjoye, A. Iqbal, K. Hadiuzzaman, X. Lu, E. Wetter, A.J. Tatem, and L. Bengtsson*

**Table S2B. Progress out of Poverty Index models for RS-only, CDR-only, and CDR+RS data: national, urban, and rural strata.**

AIC	NATIONAL	URBAN	RURAL
RS-only	41676	14099	27465
CDR-only	41562	14044	27421
CDR-RS	41502	14043	27382
<b>MODEL</b>			
RS-only	1 + annual precipitation + annual temperature + transport time to closest urban settlement + distance to roads + EVI + nighttime lights	1 + annual precipitation + annual temperature + transport time to closest urban settlement + elevation	1 + annual precipitation + annual temperature + transport time to closest urban settlement + distance to water + EVI + nighttime lights
CDR-only	1 + subscribers per tower + recharge average per tower + entropy of places + entropy of contacts + average outgoing text count + sum outgoing multimedia messages + fraction of 10 Thaka top-ups (min amount) + outgoing internet sessions + sum outgoing internet sessions + sum count incoming content management system + number of retailers visited weekly (CV) + sum revenue outgoing multimedia messages + spending speed variance + count incoming multimedia messages + sum count incoming multimedia messages + count incoming texts + sum count incoming texts + weekly recharge frequency (CV) + median time between refills + incoming video count + outgoing internet volume + time variable (CV)	1 + subscribers per tower + recharge average per tower + sum outgoing multimedia messages + count outgoing internet sessions + sum count outgoing internet sessions + number of retailers visited weekly (CV) + volume of incoming multimedia messages + outgoing text charges + sum revenue outgoing multimedia messages + incoming video duration + sum incoming video duration + spending speed variance + sum count incoming multimedia messages + weekly recharge amount + incoming video count + sum incoming video count + number of retailers visited weekly	1 + subscribers per tower + recharge average per tower + number of places + entropy of places + sum duration outgoing value added services + count outgoing texts + sum count outgoing texts + sum volume of outgoing multimedia messaging + fraction of 300 Thaka top-ups + count outgoing internet sessions + sum count outgoing internet sessions + incoming voice duration + number of retailers visited weekly (CV) + volume of incoming multimedia messages + outgoing text charges + sum outgoing text charges + sum revenue outgoing multimedia messages + spending speed variance + count incoming texts + sum count incoming texts + weekly recharge frequency (CV) + median time between refills + outgoing internet volume + recharge amount per transaction (CV) + time variable (CV)
CDR-RS	1 + annual precipitation + annual temperature + transport time to closest urban settlement + distance to road + elevation + subscribers per tower + recharge average per tower + entropy of places + entropy of contacts + average outgoing text count + sum outgoing multimedia messages + outgoing internet sessions + sum outgoing internet sessions + number of retailers visited weekly (CV) + sum revenue outgoing multimedia messages + spending speed variance + sum count incoming multimedia messages + count incoming texts + sum count incoming texts + weekly recharge frequency (CV) + median time between refills + incoming video count + outgoing internet volume + time variable (CV)	1 + annual precipitation + annual temperature + EVI + elevation + subscribers per tower + recharge average per tower + sum outgoing multimedia messages + number of retailers visited weekly (CV) + volume of incoming multimedia messages + incoming video duration + sum incoming video duration + sum revenue outgoing multimedia messages + spending speed variance + sum count incoming multimedia messages + weekly recharge amount + incoming video count + sum incoming video count + number of retailers visited weekly	1 + annual precipitation + annual temperature + transport time to closest urban settlement + distance to water + EVI + nighttime lights + subscribers per tower + recharge average per tower + outgoing multimedia messages + fraction of 300 Thaka top-ups + number of retailers visited weekly (CV) + volume of incoming multimedia messages + outgoing text charges + sum revenue outgoing multimedia messages + spending speed variance + sum revenue + median time between refills + time variable (CV)

\*CV=coefficient of variation

Table S2C. Income models for RS-only, CDR-only, and CDR+RS data: national, urban, and rural strata.

AIC	NATIONAL	URBAN	RURAL
RS-only	83194	17295	64569
CDR-only	83109	17180	64413
CDR+RS	82895	17129	64330
<b>MODEL</b>			
RS-only	1 + nighttime lights + transport time to closest urban settlement + annual temperature + EVI + distance to road + distance to water + annual precipitation + elevation	1 + nighttime lights + transport time to closest urban settlement + distance to roads	1 + nighttime lights + annual temperature + distance to roads + annual precipitation + elevation
CDR-only	1 + percent nocturnal calls + number of places + entropy of places + entropy of contacts + radius of gyration + interactions per contact + spending speed + outgoing video duration + sum outgoing video duration + average outgoing text count + sum outgoing text count + fraction of 300 Thaka top-ups + fraction of 10 Thaka top-ups + recharge amount (CV) + number of retailers visited weekly (CV) + recharge amount per transaction + spending speed variance + sum spending speed variance + sum count incoming texts + handset weight + outgoing voice duration + sum outgoing voice duration + sum outgoing internet volume + time variable (CV) + fraction of 200 Thaka top-ups	1 + percent nocturnal calls + number of places + entropy of places + radius of gyration + spending speed + outgoing video duration + sum outgoing video duration + average outgoing text count + sum outgoing text count + fraction of 300 Thaka top-ups + number of retailers visited weekly (CV) + volume of incoming multimedia messages + sum volume of incoming multimedia messages + recharge amount per transaction + count incoming multimedia messages + sum count incoming multimedia messages + sum outgoing voice duration + outgoing internet volume + sum outgoing internet volume + recharge amount per transaction (CV) + time variable (CV)	1 + number of places + entropy of places + entropy of contacts + spending speed + duration outgoing value added services + sum duration outgoing value added services + sum outgoing video duration + handset weight + software OS version + fraction of 300 Thaka top-ups + fraction of 10 Thaka top-ups + coefficient of variation: recharge amount + number of retailers visited weekly (CV) + volume of incoming multimedia messages + sum volume of incoming multimedia messages + spending speed variance + sum spending speed variance + sum count incoming multimedia messages + count incoming texts + sum count incoming texts + weekly recharge amount + outgoing voice duration + sum outgoing voice duration + time variable (CV)
CDR-RS	1 + nighttime lights + transport time to closest urban settlement + EVI + distance to road + percent nocturnal calls + number of places + entropy of places + spending speed + outgoing video duration + sum outgoing video duration + average outgoing text count + sum outgoing text count + recharge amount (CV) + number of retailers visited weekly (CV) + recharge amount per transaction + spending speed variance + sum spending speed variance + sum count incoming texts + handset weight + outgoing voice duration + sum outgoing voice duration + sum outgoing internet volume + time variable (CV) + fraction of 200 Thaka top-ups	1 + nighttime lights + transport time to closest urban settlement + annual temperature + distance to roads + radius of gyration + spending speed + outgoing video duration + sum outgoing video duration + average outgoing text count + sum outgoing text count + fraction of 300 Thaka top-ups + number of retailers visited weekly (CV) + volume of incoming multimedia messages + recharge amount per transaction + count incoming multimedia messages + sum count incoming multimedia messages + sum outgoing voice duration + outgoing internet volume + sum outgoing internet volume + recharge amount per transaction (CV) + time variable (CV)	1 + nighttime lights + annual precipitation + number of places + entropy of places + entropy of contacts + spending speed + sum outgoing video duration + handset weight + software OS version + fraction of 300 Thaka top-ups + fraction of 10 Thaka top-ups + coefficient of variation: recharge amount + number of retailers visited weekly (CV) + weekly recharge amount + outgoing voice duration + sum outgoing voice duration + outgoing internet volume + time variable (CV)

\*CV=coefficient of variation

Table S3. Comparison of  $r^2$  and RMSE for INLA models run with only a structured spatial random effect (Spatial interpolation) and the full model (Spatial model + covariates).

Poverty Metric	Spatial interpolation	Spatial model + covariates (from Table 1)
	$R^2$ , RMSE	$R^2$ , RMSE
DHS WI	0.49, 0.578	0.76, 0.394
PPI	0.31, 58.727	0.32, 57.439
Income	0.10, 123.963	0.27, 105.465

Table S4. Comparison of deviance information criterion (DIC) model fit for Bayesian geostatistical models run with a structured spatial random effect (Spatial model) and without (Non-spatial model).

<b>WHOLE COUNTRY</b>			
Poverty Metric	Model	Spatial model	Non-spatial model
		DIC	DIC
DHS WI	CDR -		
	RS	463.7	574.6
	CDR	272.5	862.2
PPI	RS	465.7	581.5
	CDR -		
	RS	1361.3	1439.8
Income	CDR	1349.7	1473.1
	RS	1358.6	1421.4
	CDR -		
	RS	66142.5	66143.0
	CDR	66314.6	66314.3
	RS	66482.6	66480.4
<b>URBAN</b>			
Poverty Metric	Model	Spatial model	Non-spatial model
		DIC	DIC
DHS WI	CDR -		
	RS	449.4	576.0
	CDR	239.2	873.2
PPI	RS	454.1	582.0
	CDR -		
	RS	1371.6	1432.3
Income	CDR	1365.7	1470.1
	RS	1358.3	1417.7
	CDR -		
	RS	66180.6	66179.8
	CDR	66363.3	66365.7
	RS	66693.0	66690.3
<b>RURAL</b>			
Poverty Metric	Model	Spatial model	Non-spatial model
		DIC	DIC
DHS WI	CDR -		
	RS	458.0	574.5
	CDR	63.1	873.5
PPI	RS	451.5	595.5
	CDR -		
	RS	1376.9	1444.6
Income	CDR	1342.9	1475.9
	RS	1357.7	1419.4
	CDR -		
	RS	66262.5	66260.6
	CDR	66395.0	66392.1
	RS	65548.9	66503.8

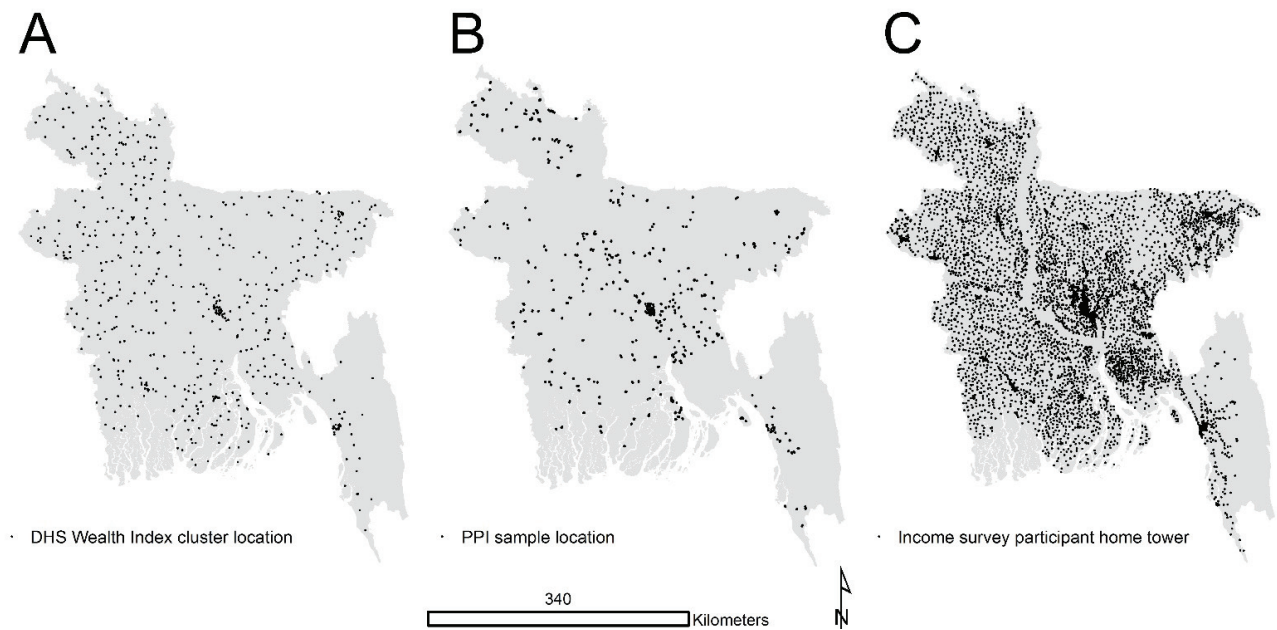


Figure S1. Survey sample locations for DHS wealth index (A), Progress out of Poverty Index (B), and income survey respondents (C).



# Mapping poverty using mobile phone and satellite data

*J.E. Steele, P. Sundsoy, C. Pezzulo, V. Alegana, T. Bird, J. Blumenstock, J. Bjelland, YA de Montjoye, K. Engo-Monsen, A. Iqbal, K. Hadiuzzaman, X. Lu, E. Wetter, L. Bengtsson, and A.J. Tatem*

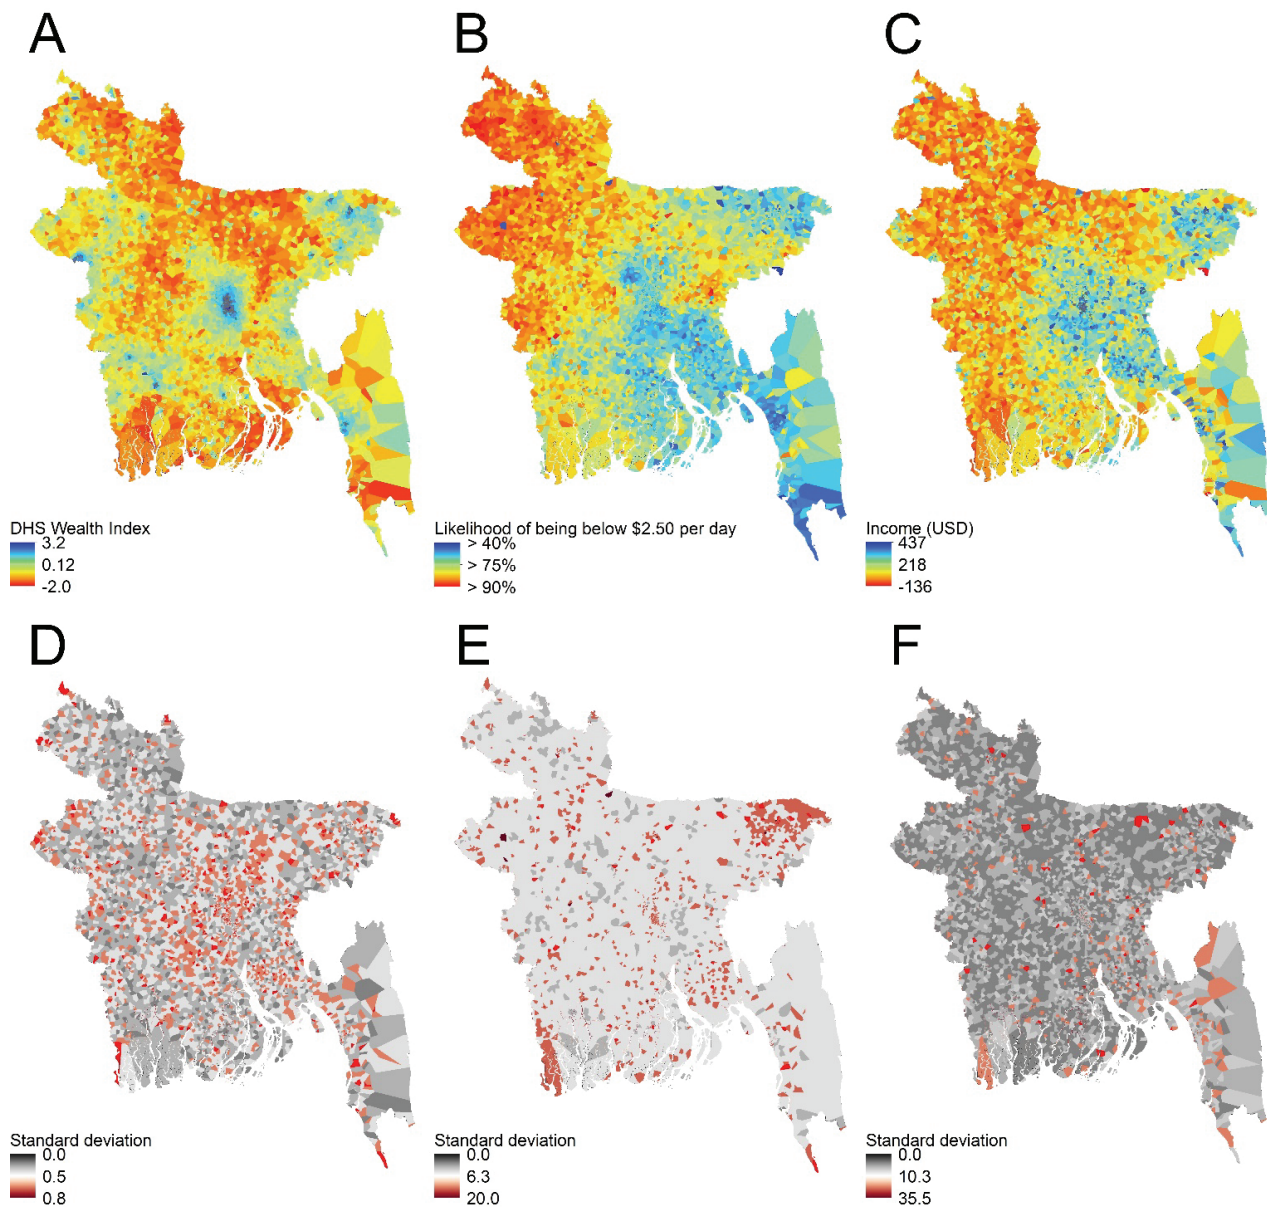


Figure S2. National level prediction maps for mean wealth index (A) with uncertainty (D); mean probability of households being below \$2.50/day (B) with uncertainty (E); and mean USD income (C) with uncertainty (F). Maps were generated using call detail record features only and Bayesian geostatistical models. Red indicates poorer areas in prediction maps, and higher error in uncertainty maps.

# Mapping poverty using mobile phone and satellite data

*J.E. Steele, P. Sundsoy, C. Pezzulo, V. Alegana, T. Bird, J. Blumenstock, J. Bjelland, YA de Montjoye, K. Engo-Monsen, A. Iqbal, K. Hadiuzzaman, X. Lu, E. Wetter, L. Bengtsson, and A.J. Tatem*

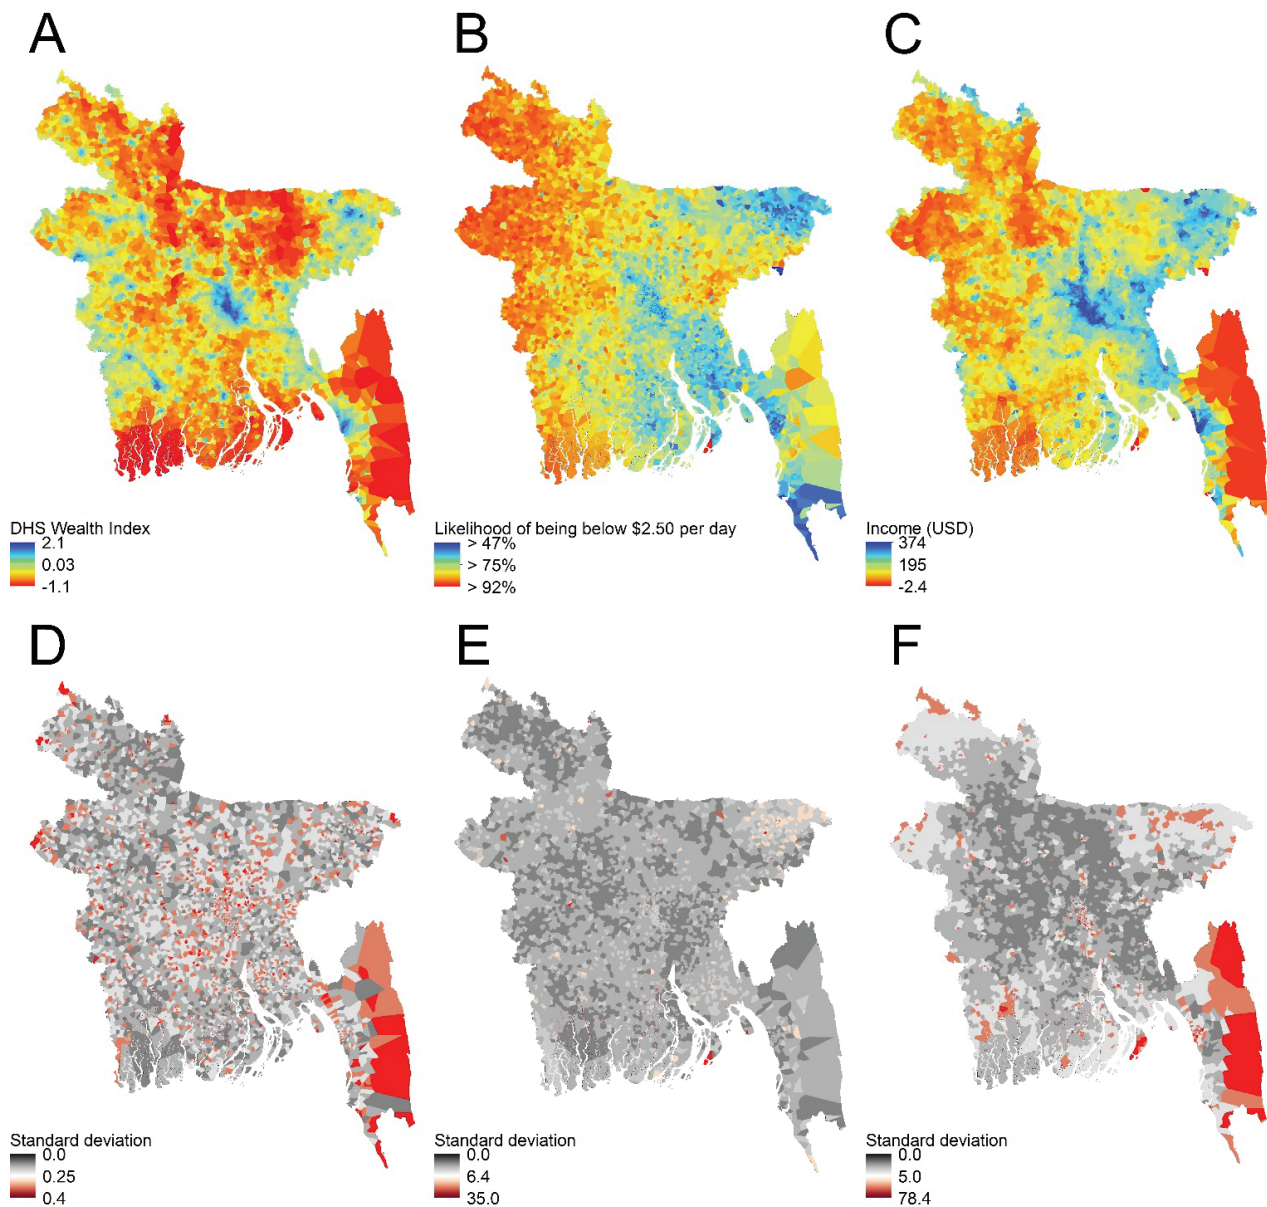


Figure S3. National level prediction maps for mean wealth index (A) with uncertainty (D); mean probability of households being below \$2.50/day (B) with uncertainty (E); and mean USD income (C) with uncertainty (F). Wealth index and income maps were generated using remote sensing data only; PPI maps were generated using call detail record features and remote sensing data. All maps were generated using Bayesian geostatistical models. Red indicates poorer areas in prediction maps, and higher error in uncertainty maps.



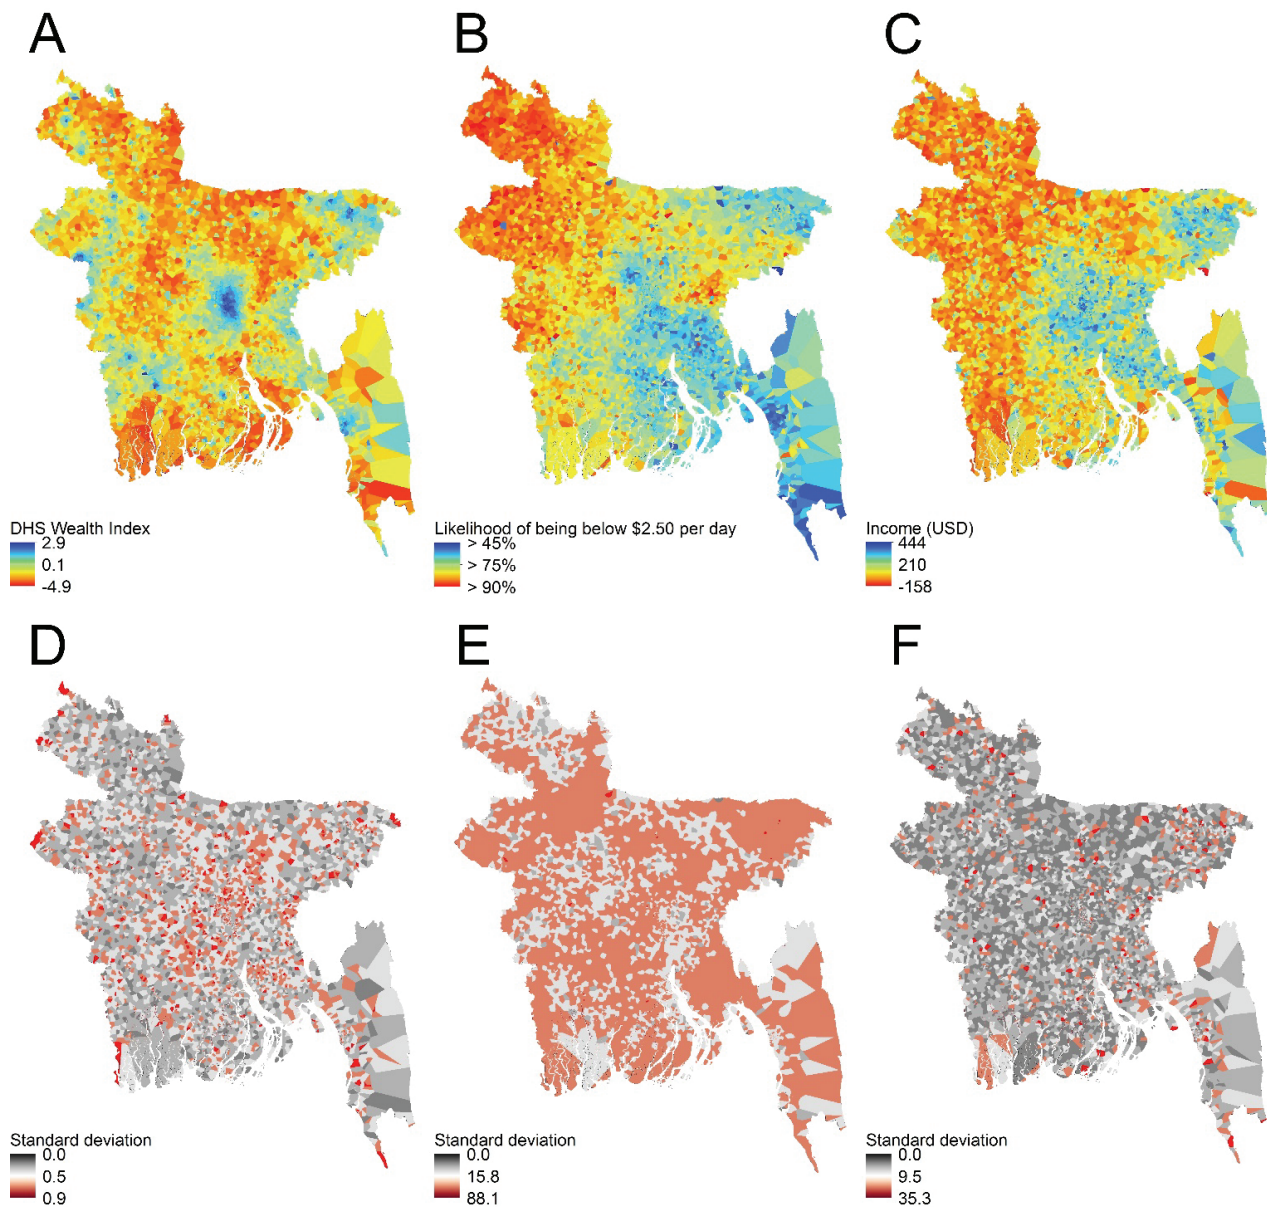


Figure S4. Stratified urban/rural prediction maps for mean wealth index (A) with uncertainty (D); mean probability of households being below \$2.50/day (B) with uncertainty (E); and mean USD income (C) with uncertainty (F). Maps were generated using call detail record features only and Bayesian geostatistical models. Red indicates poorer areas in prediction maps, and higher error in uncertainty maps.

# Mapping poverty using mobile phone and satellite data

*J.E. Steele, P. Sundsoy, C. Pezzulo, V. Alegana, T. Bird, J. Blumenstock, J. Bjelland, YA de Montjoye, K. Engo-Monsen, A. Iqbal, K. Hadiuzzaman, X. Lu, E. Wetter, L. Bengtsson, and A.J. Tatem*

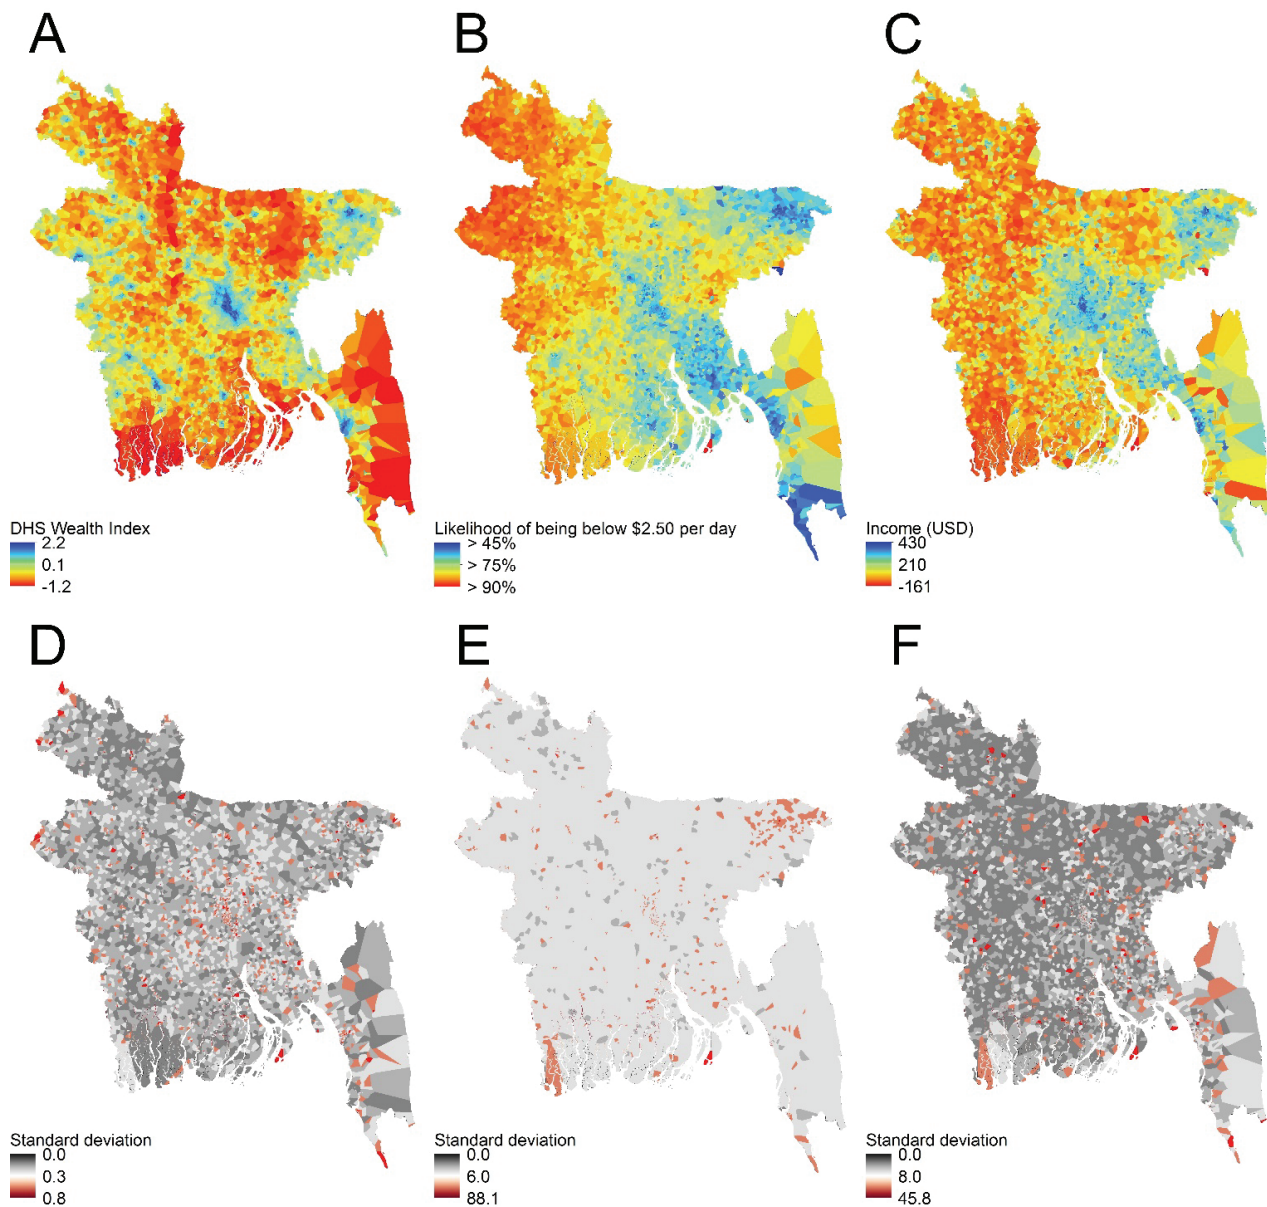


Figure S5. Stratified urban/rural prediction maps for mean wealth index (A) with uncertainty (D); mean probability of households being below \$2.50/day (B) with uncertainty (E); and mean USD income (C) with uncertainty (F). Maps were generated using call detail record features, remote sensing data, and Bayesian geostatistical models. Red indicates poorer areas in prediction maps, and higher error in uncertainty maps.



# Mapping poverty using mobile phone and satellite data

*J.E. Steele, P. Sundsoy, C. Pezzulo, V. Alegana, T. Bird, J. Blumenstock, J. Bjelland, YA de Montjoye, K. Engo-Monsen, A. Iqbal, K. Hadiuzzaman, X. Lu, E. Wetter, L. Bengtsson, and A.J. Tatem*

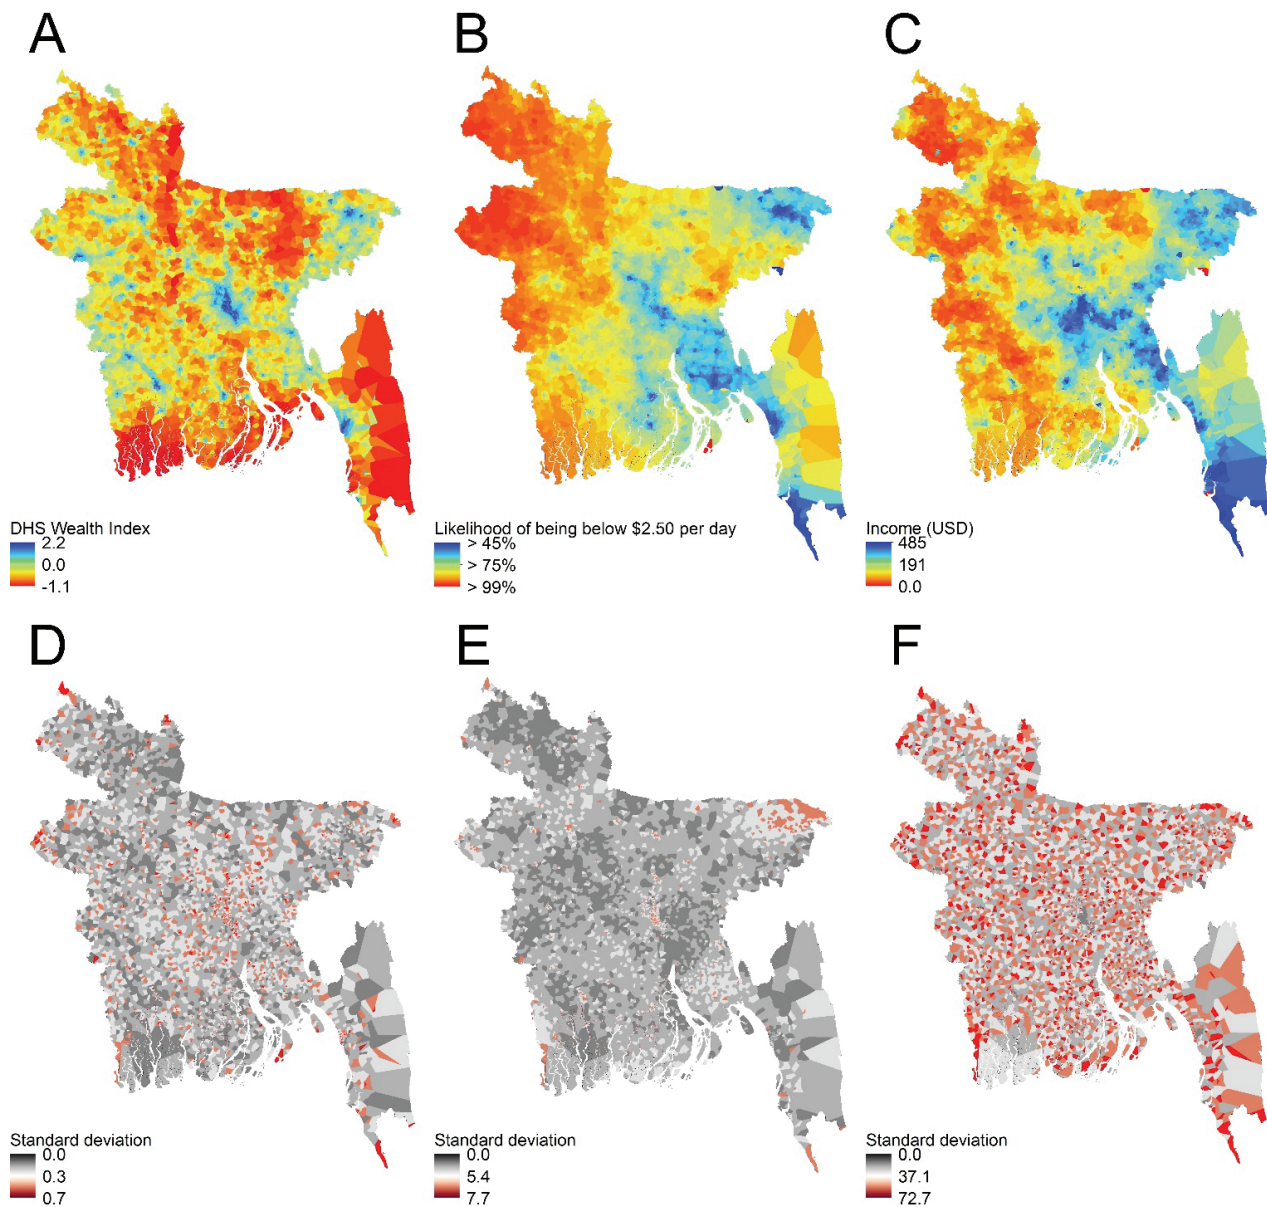


Figure S6. Stratified urban/rural prediction maps for mean wealth index (A) with uncertainty (D); mean probability of households being below \$2.50/day (B) with uncertainty (E); and mean USD income (C) with uncertainty (F). Maps were generated using remote sensing data only and Bayesian geostatistical models. Red indicates poorer areas in prediction maps, and higher error in uncertainty maps.



## **Paper 4**

---



## **Paper 5**

---





# Detecting climate adaptation with mobile network data in Bangladesh: anomalies in communication, mobility and consumption patterns during cyclone Mahasen

Xin Lu<sup>1,2,3</sup> · David J. Wrathall<sup>4</sup> · Pål Roe Sundsøy<sup>5</sup> ·  
Md. Nadiruzzaman<sup>6,7</sup> · Erik Wetter<sup>2,8</sup> · Asif Iqbal<sup>5</sup> ·  
Taimur Qureshi<sup>5</sup> · Andrew J. Tatem<sup>2,9</sup> ·  
Geoffrey S. Canright<sup>5</sup> · Kenth Engø-Monsen<sup>5</sup> ·  
Linus Bengtsson<sup>1,2</sup>

Received: 9 February 2016 / Accepted: 8 July 2016 / Published online: 1 August 2016  
© The Author(s) 2016. This article is published with open access at Springerlink.com

**Abstract** Large-scale data from digital infrastructure, like mobile phone networks, provides rich information on the behavior of millions of people in areas affected by climate stress. Using anonymized data on mobility and calling behavior from 5.1 million Grameenphone users in Barisal Division and Chittagong District, Bangladesh, we investigate the effect of Cyclone Mahasen, which struck Barisal and Chittagong in May 2013. We characterize spatiotemporal patterns and anomalies in calling frequency, mobile recharges, and population movements before, during and after the cyclone. While it was originally anticipated that the analysis might

---

**Electronic supplementary material** The online version of this article (doi:10.1007/s10584-016-1753-7) contains supplementary material, which is available to authorized users.

---

✉ Linus Bengtsson  
linus.bengtsson@flowminder.org

<sup>1</sup> Department of Public Health Sciences, Karolinska Institutet, Stockholm, Sweden

<sup>2</sup> Flowminder Foundation, Stockholm, Sweden

<sup>3</sup> College of Information System and Management, National University of Defense Technology, Changsha, China

<sup>4</sup> College of Earth, Ocean and Atmospheric Sciences, Oregon State University, Corvallis, OR, USA

<sup>5</sup> Telenor Research, Oslo, Norway

<sup>6</sup> Department of Geography, University of Exeter, Exeter, UK

<sup>7</sup> International Centre for Climate Change and Development, Dhaka, Bangladesh

<sup>8</sup> Stockholm School of Economics, Stockholm, Sweden

<sup>9</sup> WorldPop, Department of Geography and Environment, University of Southampton, Southampton, UK

detect mass evacuations and displacement from coastal areas in the weeks following the storm, no evidence was found to suggest any permanent changes in population distributions. We detect anomalous patterns of mobility both around the time of early warning messages and the storm's landfall, showing where and when mobility occurred as well as its characteristics. We find that anomalous patterns of mobility and calling frequency correlate with rainfall intensity ( $r = .75$ ,  $p < 0.05$ ) and use calling frequency to construct a spatiotemporal distribution of cyclone impact as the storm moves across the affected region. Likewise, from mobile recharge purchases we show the spatiotemporal patterns in people's preparation for the storm in vulnerable areas. In addition to demonstrating how anomaly detection can be useful for modeling human adaptation to climate extremes, we also identify several promising avenues for future improvement of disaster planning and response activities.

**Keywords** Climate change adaptation · Migration · Resilience · Mobile network data · Anomaly detection · Disaster risk

## 1 Introduction

The increasingly robust evidence base in climate sciences relies on the measurement of normal trends and analysis of deviations (Bindoff et al. 2013). Techniques for detecting anomalies have produced key findings on changing atmospheric and surface temperature (Jones et al. 1999; Mann et al. 1998), oceanic circulation (Hurrell 1995; Thompson and Wallace 1998), arctic temperatures and ice cover (Serreze et al. 2000; Stroeve et al. 2007; Vinje 2001), intensity of tropical rainfall and cyclones (Knutson et al. 2010; Trenberth 2011); and seasonal variability and extremes (Seneviratne et al. 2012). Anomaly detection principles have also shown how earth's ecosystems (Lucht et al. 2002; Stenseth et al. 2002) and biota including agriculture (Lenoir et al. 2008) have responded to climate change. However, human behavior in response to disasters also deviates from normal behavioral patterns. In this paper, we aim to use anomaly detection to investigate behavioral responses in a human population exposed to an extreme weather event.

Vulnerable people in low- and middle-income countries respond to weather extremes associated with climate change, such as tropical cyclones and flooding, with a variety of behaviors that appear anomalous against a baseline (here termed “adaptations”) such as moving animals to safety, harvesting crops early, reinforcing and repairing flood embankments, and changing household spending behaviors. In more extreme cases, short-term adaptive responses include evacuation and displacement. Weather extremes can, in the long-term, undermine livelihoods, push people into poverty, and elicit an extraordinary adaptive response in these circumstances: permanent migration (Black et al. 2011; Brouwer et al. 2007), a subject of rich academic debate (summarized in (Black et al. 2013) featured centrally in 5th Assessment of the IPCC (Adger et al. 2014; Olsson et al. 2014). Unfortunately, our ability to detect anomalous human behaviors is not on par with our large-scale measurements of biophysical systems at relevant temporal and spatial scales.

Climate science has seen rapid progress in the measurement, and prediction of changes and extremes in biophysical systems in high resolution across geographic and temporal scales. To understand the impacts of climate change on human society it is imperative to measure anomalous behavioral responses as they coincide with hazards at the common spatiotemporal scales in which they occur (Palmer and Smith 2014). This is especially crucial where people are dependent on stable environmental conditions for livelihoods, and where both climate change and the burden of

adaptation threaten human security and development (Adger et al. 2014; Field et al. 2014). Methodologies that focus on large-scale spatial indicators of both human behavioral and environmental change, and make use of temporally adjusted longitudinal data are required to establish baselines and link short-term responses and long-term outcomes (Palmer and Smith 2014).

As of the end of 2014, mobile networks served a total of 3.6 billion unique mobile subscribers, roughly half of the global population (GSMA Intelligence 2015). Mobile operator data are updated in close to real-time and have a vast geographic reach. The data generated from mobile operators enable measurement of some characteristics of social networks, migration, and patterns of household economic behavior at a previously unprecedented scale (Bagrow et al. 2011; Palmer and Smith 2014; Zolli 2012). Operator data has been used during relief operations after the Haiti 2010 earthquake (Bengtsson et al. 2011; Lu et al. 2012) and cholera outbreaks (Bengtsson et al. 2015) and the Nepal 2015 earthquake (Wilson et al. 2016), making them a very promising proxy indicator for measuring impacts of climate change, and weather extremes. In Rwanda, retrospective analyses of network activity was used to estimate the epicenter of an earthquake and to infer humanitarian needs in the weeks after the earthquake (Kapoor et al. 2010). Likewise, Blumenstock and colleagues identified unusual patterns of person-to-person transfers of airtime credits through social networks to identify a geographical pattern of earthquake impact (Blumenstock et al. 2011). Anomaly detection methods have previously been applied to mobile network data to identify unusual calling patterns after floods (Pastor-Escuredo et al. 2014), and in the interest of improving the normal operation of mobile networks (Karatepe and Zeydan 2014). They have been used for anomaly detection for detecting and classifying social disturbances, like conflict and violence in data-poor circumstances (Dobra et al. 2014; Young et al. 2014). One study showed the diffusion of anomalous calling patterns through intimate social networks in the wake of a terrorist bombing in Oslo (Sundøy et al. 2012). Various studies have concluded that in the wake of disasters anomaly detection could reduce the cost, increase timeliness and improve the geographic focus of emergency response activities (Candia et al. 2008; Pawling et al. 2007).

The extreme South of coastal Bangladesh, with its low elevation and routine exposure to intense tropical cyclones, exemplifies an area with high climate pressure and is a fitting location to explore mobile network data before and after climate shocks. We searched for anomalous patterns of phone usage that could provide insight into adaptive preparations and responses (Martin et al. 2014; McGranahan et al. 2007; Penning-Rowsell et al. 2013), and examined how spatial and temporal patterns in large sets of operator data from the Grameenphone mobile network in Bangladesh around tropical cyclone Mahasen could inform impact assessment and adaptation in cyclone affected areas. We investigated three hypotheses. First, anomalous patterns of calling frequency represent the affected populations' physical contact with the storm in the most affected areas during landfall. Second, as communication is an important tool during an environmental crisis, we hypothesized that anomalous mobile recharges represent behaviors of people preparing for impacts in the most vulnerable areas. Finally, we hypothesized that cyclones drive anomalous flows of users between towers, indicating evacuation, displacement and migration.

## 2 Cyclone Mahasen

Cyclone Mahasen struck Bangladesh on 16 May 2013. Before landfall it moved northward along the Bay of Bengal. Forecasts estimated a landfall in the heavily populated Chittagong

District, and the government's Comprehensive Disaster Management Programme concentrated early warnings there. However, in the final hours of 15 May, the storm veered to the north, making landfall over the rural Barisal Division, at approximately 3:00 a.m. on 16 May (Fig. S1a). During the course of 16 May, Mahasen moved eastward along the coast into Chittagong, and northward into India, where rainfall and wind speed rapidly diminished (Gutro and Pierce 2013).

Mahasen was a relatively weak storm compared to earlier cyclones in Bangladesh, such as Aila and Sidr (REACH Initiative 2013). While it affected an estimated 1.3 million people (REACH Initiative 2013) and impacts on crops and homes were extensive, the death toll was relatively small. Seventeen perished in the storm, mostly from falling trees, and unlike previous storms, no fishermen were lost (Associated Press 2013). The minimal loss of life was regarded as a major victory for the Comprehensive Disaster Management Programme's early warning system (UNDP 2013).

### 3 Cyclone impacts and population-level adaptation

#### 3.1 Mobile phone data

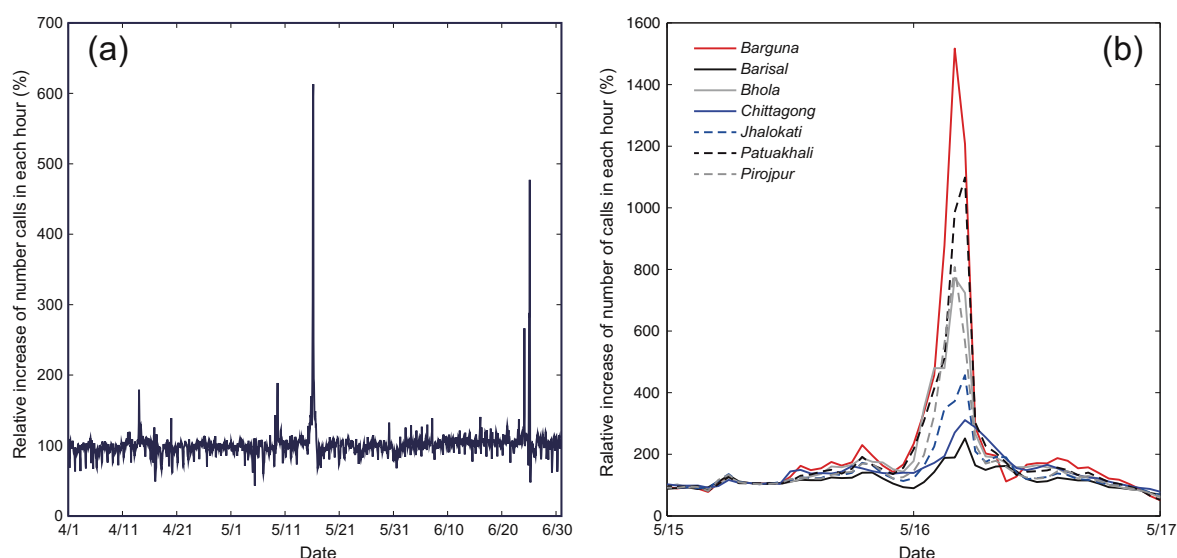
We used a de-identified set of call detail records (CDRs) from 5.1 million Grameenphone users collected between 1 April and 30 June of 2013 in the Barisal Division and Chittagong District of Bangladesh. The dataset began six weeks before the landfall of Cyclone Mahasen (16 May 2013) and continued for six weeks after landfall (1 April to 30 June 2013) (Fig. S1). CDRs are compiled by network operators principally for the purposes of billing customers for their use of the network. De-identified data entries include information on the time of the call, the mobile phone tower used and the duration of call, and can thus be used to indicate the geographical position and movements of users. To limit potential biases resulting from subscriber churn, and new users entering the dataset due to impacts of the storm, we limited the study to SIM cards that had placed at least one call before the cyclone landfall (16 May); and also made at least one call in the last ten days of the data collection period (21–30 June).

Since Mahasen was a relatively weak cyclone, the performance of the Grameenphone network remained virtually undisturbed during and after landfall, guaranteeing continuous relay of CDR throughout the study period. An analysis of tower function anomalies appears in the *Supporting Information* (S2), along with a general discussion on the Grameenphone network, the dataset, and the representativeness of data for the general population (S1).

#### 3.2 Calling frequency and rainfall measurements

During “normal” circumstances, which we defined as the average calls per hour for any given hour across the data set, regular daily and weekly cycles of calls were apparent (Fig. 1a). Users concentrated phone usage in the daytime hours, with a spike occurring toward late evening. A small shift in the temporal distribution of calls occurred on Fridays (the first day of the weekend) when calls began later in the day. Increases in calling frequency coincided with several events within the data set, most notably on 25 June, when a major religious festival, Shab-e-Barat was celebrated. Likewise smaller increases coincided with the Bengali New Year in early April and a series of protests in early May.

However in the early hours of Thursday 16 May 2013, as Mahasen made landfall across Barisal, we observed a dramatic increase in call frequency relative to “normal,” which we defined



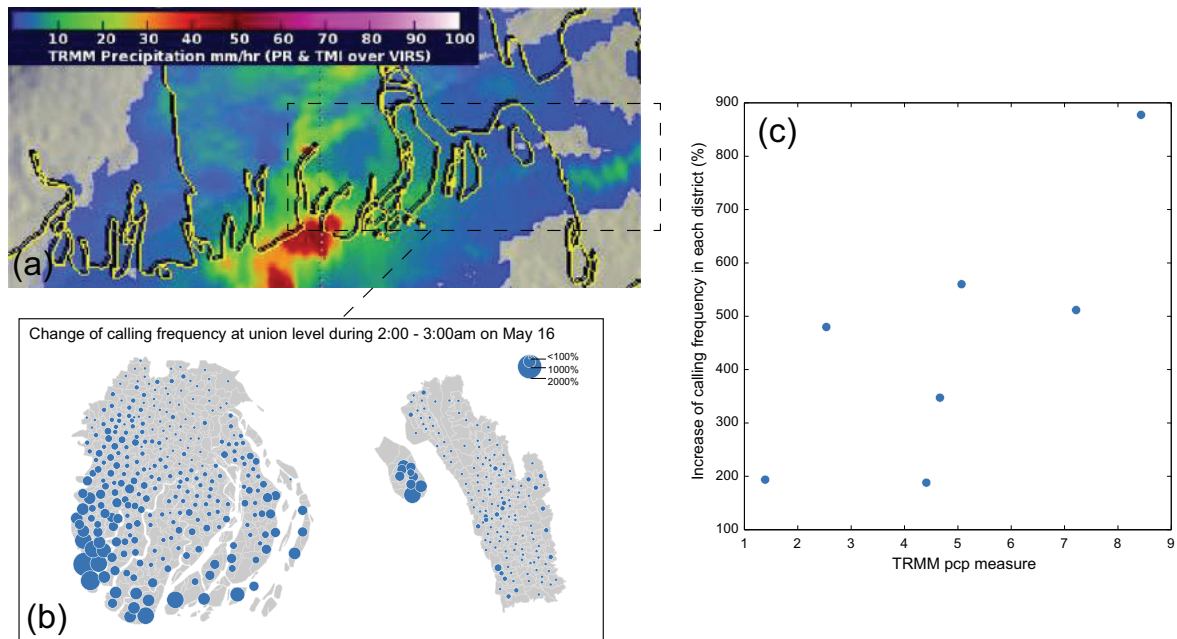
**Fig. 1** Change in call frequency. **a** For each hour, the number of calls is compared to the average number of calls made during that same hour across the whole period. Relative calling frequency spikes on 16 May as Mahasen makes landfall. **b** Calling frequency during cyclone landfall at the district level. For each of the seven districts, the change of calling frequency in the morning of 16 May is shown

as calls per hour compared with the same hour on all Thursdays in the dataset (Fig. 1a). Deconstructing calling frequency among the more vulnerable coastal districts (Barguna, Bhola, Patuakhali, Pirojpur), we saw calling volumes increase by at least seven times the average level (Fig. 1b). In Barguna, calling frequency increased by a factor of 15. Throughout the evening and early hours, a spatiotemporal pattern emerged in peak calling frequency. In Barguna and Pirojpur, in the extreme south and west of Barisal, peaks occurred between 3:00 and 4:00 a.m., while in the northern districts of Jhalokhati and Barisal, the peak occurred between 4:00 and 5:00 a.m. This suggests that call frequency corresponded with the physical manifestations of the cyclone as it moved over Barisal Division from the South West. One alternative explanation is that people might call friends and relatives when a cyclone is approaching in order to communicate concerns for wellbeing, encourage evacuation plans, and coordinate preparations. However, between 00:00 and 6:00 a.m. at the time Mahasen was making landfall in Barisal, in the areas where the cyclone was predicted to make landfall (the Chittagong district), calling frequency was close to normal levels. These differences provide support for the hypothesis that calling frequency represented a behavioral response to sensory experience of the storm.

To further investigate the relationship between calling frequency and physical manifestations of the storm, we conducted a spatiotemporal comparison of calling frequency with rainfall data from NASA's Tropical Rainfall Measurement Mission (TRMM) satellite. The TRMM satellite passed over Bangladesh at 3:32 a.m., measuring rainfall during the cyclone's landfall, reaching 67 mm per hour in some areas (Gutro and Pierce 2013). Locations of maximum rainfall were clearly correlated with locations of maximum increase in calling frequency (Fig. 2). Even areas with moderate rainfall, for example a narrow band of rainfall to the east of Chittagong (Fig. 2a) also exhibited an increase in calling frequency (Fig. 2b). This adds supporting evidence that clusters of high calling frequency represented contact with the cyclone's most severe physical effects.

Past research has shown that rainfall data alone is often too low resolution and intermittent to make any inferences about cyclone damage (Auffhammer et al. 2013). Detailed spatiotemporal data on call frequency may improve inferences about the effect of weather extremes on vulnerable people, and is identified here as an area for future research.

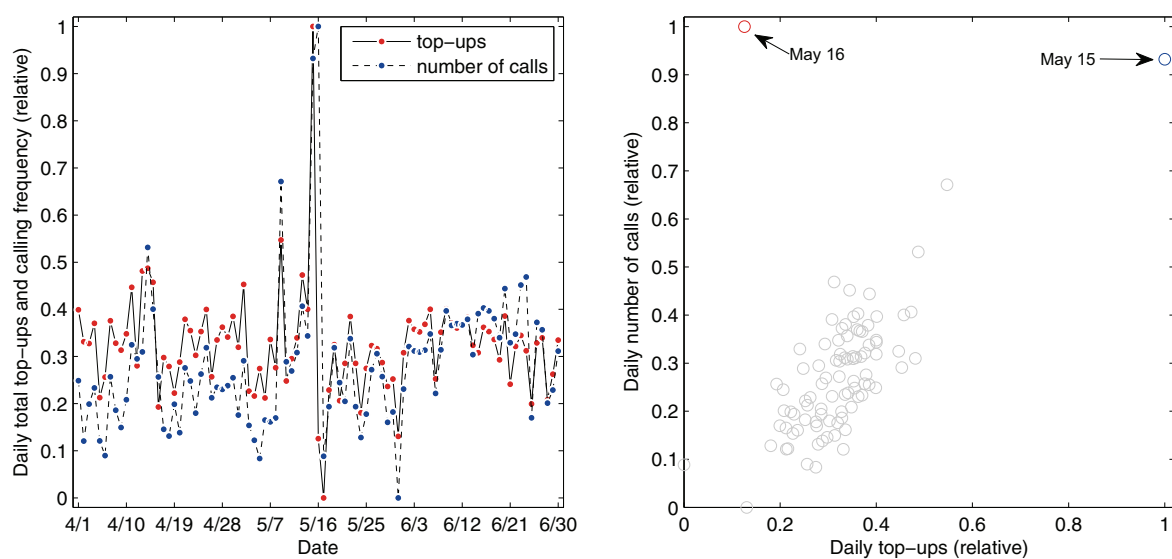




**Fig. 2** Call anomalies and rainfall. **a** Precipitation measurements from NASA's Tropical Rainfall Measurement Mission captured at 3:32 a.m. show the distribution of rainfall in the study area, reprinted from Gutro and Pierce 2013, with permission from the authors. **b** The geographical distribution of call frequency at 3:00 a.m. on 16 May. **c** Rainfall is plotted with calling frequency at the district level. Correlation coefficient = 0.75,  $p = 0.05$

### 3.3 Recharge behaviors

Next, we investigated how mobile recharges or top-ups can complement call frequency to provide insight on how vulnerable people prepare for climate impacts. To accomplish this, we relied on a second data set, consisting of mobile recharge purchases from 892 retailers in Barisal and Chittagong Divisions during the original three-month timeline, 1 April to 30 June of 2013 (Fig. 3). Recharges are the amount of money that users credit to their SIM card to



**Fig. 3** Comparison of daily recharges and calling frequency, both features are presented with min-max normalization. **a** During the cyclone, while calling frequency is high for both 15 May and 16 May, recharges are high only on 15 May and then drop to low levels on 16 May. **b** Excluding 16 May, there is a strong linear relationship between recharges and calling frequency (corr without 16 May = 0.798,  $p < 0.000$ , corr with 16 May = 0.576,  $p < 0.000$ ), implying that recharges indicate users preparing for a potential disaster



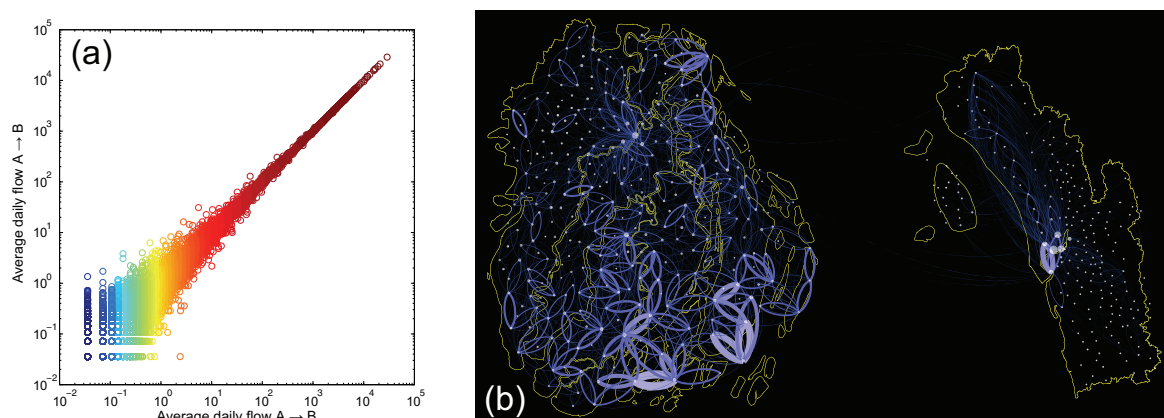
access the network. They allowed an investigation of the geographic distribution of changes in expenditures before and after the cyclone. In Bangladesh, mobile credits represent a small but significant proportion (~3 %) of the household budget (Lucini and Hatt 2014), and disasters increase demand for private communication (Samarajiva 2005). We hypothesized that spikes in recharges represented knowledge of the cyclone and preparations for its impacts.

In the second half of 15 May, as forecasts and early warnings were transmitted across radio and television, a large increase in recharges is evident, coinciding with a high volume of calls placed on the same day (Fig. 3). However even as calling frequency remained high on 16 May, recharges dropped below the predicted level, and continued at low levels during the following day. This suggests that users recharged their phones as part of their storm preparation and awareness of vulnerability, planning for the need to communicate with family and friends during and after the cyclone.


### 3.4 Estimating evacuation, displacement and migration

Usage patterns in the data also enabled us to analyse short-term features of evacuation, displacement and migration, which would be extremely hard to quantify using standard survey-based research but were readily apparent in CDRs. Using CDRs and tower locations to identify moving SIM cards, we created a series of mobility networks, which quantify the direction, volume and distance of flows between locations at specified time intervals before and after Cyclone Mahasen (Figs. 4 and 5). Note that the mobility networks during normal periods were almost perfectly symmetrical, meaning that the numbers of users entering an area are roughly equal to the number of users leaving an area (Fig. 4). In contrast, anomalies appeared as larger than normal flows in one or both directions (Fig. 5), and indicated spatiotemporally explicit patterns of movement, such as evacuation, displacement and permanent migration that took place at specific moments coinciding with the storm. Because asymmetrical flows might also represent, for example, the onset of migration season, a calendar festival or a popular protest, it is important to be cautious in assigning causation.

Prior to the storm, large changes in the flow network were notable in Chittagong City, as people evacuated in response to the forecasts that Mahasen would make landfall over



**Fig. 4** Symmetry in “normal” flow networks. **a** The average daily flow between each union in both directions between April 1, to April 28 (i.e., for four “normal” weeks prior to the storm) are shown. The correlation is extremely high ( $\text{corr} = 0.999$ ,  $p < 0.000$ ), indicating a high level of symmetry in “normal” mobility patterns, i.e. each day, roughly as many people leave an area as those who enter it. **b** This symmetry of “normal” flow can be represented geospatially. SIM cards are included only if they accessed more than one tower in a day. Links indicate areas where 10 or more movements were observed at distances greater than 10 km. Direction of flow is clockwise from the point of origin



**Fig. 5** Evacuation and landfall flow networks. **a** The mobility network on 15 May (the day prior to landfall) is compared with 24 April (3 weeks before the storm during the same hourly period). Positive flows are shown in *red*, indicating increased flow on 15 May, while negative flows are shown in *blue*, indicating decreased flow on 15 May. Thickness of link represents relative volume of flow. To appear in the flow network, a user had to make at least two calls. Each SIM contributed only one movement (the first and last observed location). Links indicate areas where 10 or more movements were observed, at distances greater than 10 km. **b** The mobility network during landfall on 16 May, 00:00–6:00 a.m., is compared with 25 April (3 weeks prior during the same hourly period). Unusual mobility is observed in the affected area, where warnings were not concentrated

Chittagong (Fig. 5a) Meanwhile, there were less than normal flows in Barisal at the same time, suggesting people were not evacuating to other areas in large numbers, but rather suspending regular trips.

In the early hours of May 16, during cyclone landfall, at the time when people should have been in shelters, above normal flows of SIMs were evident in the margin of sub-districts in Barisal nearest to the coast, indicating that people were moving about at night, during the storm (Fig. 5b). This suggests that people evacuated too late, and would have been in danger if the storm intensity had been greater. Mobility patterns in Barisal during landfall contrasted sharply with mobility in Chittagong during the same time, where patterns were virtually unchanged from normal flows for that day and hour. Although the exact explanation for these differences is unknown, officials in the Ministry of Disaster Planning and Response indicated that early warnings were not made in Barisal until too late because all forecasting indicated that the storm would make landfall in Chittagong (Nadiruzzaman 2013). Other possible explanations for delayed evacuation in rural areas were that men commonly stay behind to look after livestock and protect homes and assets from thieves.

In sum, the mobility patterns evident in mobile network data allow researchers to perform an audit of early warning program effectiveness on the basis of early and mid-storm population movements. In this case, the early warning system in Chittagong apparently accomplished the aim of motivating evacuation during appropriate times.

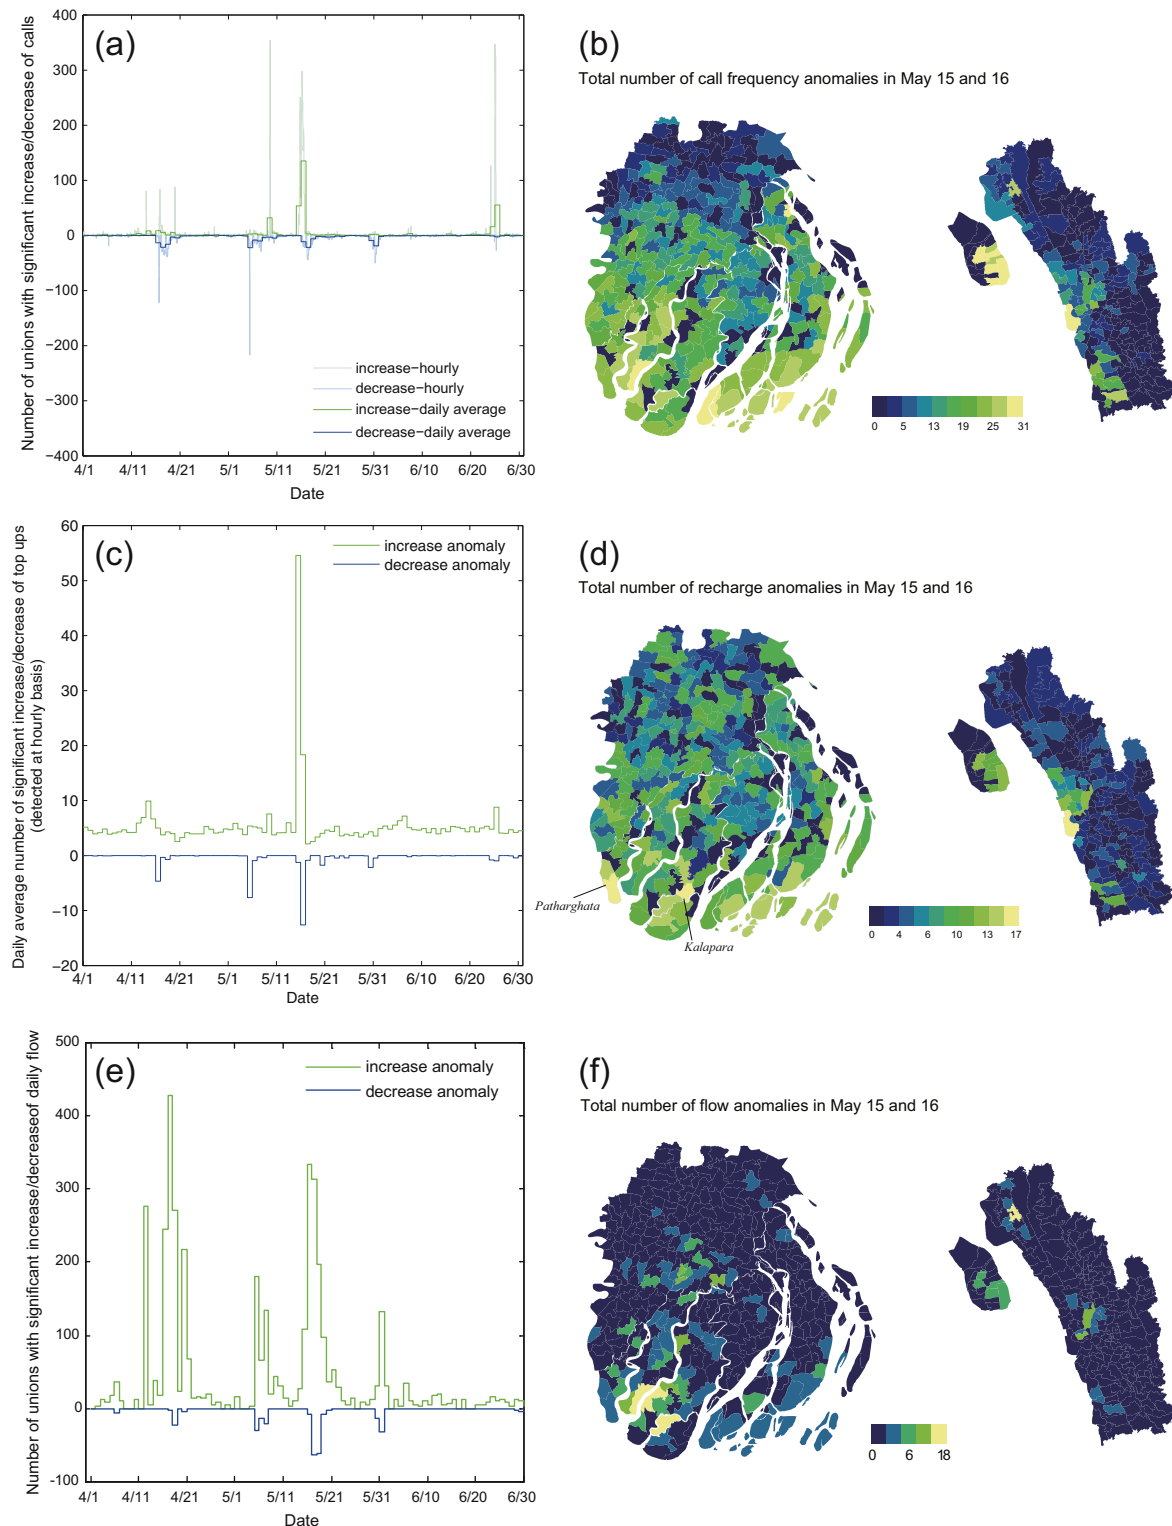
## 4 Quantifying impacts and behavioral responses using anomaly detection techniques

To automatically detect human behavioral changes in our study, we used a sigma-model to evaluate the stability of the observed sequence of activities extracted from customers' usage data in the mobile network. Specifically, for each time series of a quantified activity,  $E = \{e_1, e_2, e_3, \dots, e_t\}$ , in which  $e_i \in \mathbb{R} (1 \leq i \leq t)$  is the measure of the activity (e.g., number of calls at time  $i$ , and  $t$  the length of evaluated time window) we highlight time points  $I = \{1 \leq i_1, i_2, \dots, i_M \leq t\}$  in which each observation  $e_{i_m}$  at time point  $i_m$  exceeds three standard deviations from  $E$ 's average during the time period.

As the studied activity reached the predetermined thresholds, three standard deviations from the mean, it was flagged as anomalous. In this way, any unusual patterns of network usage could be identified, and further analysis would determine what these anomalies represented about cyclone impacts. When the number of anomalous cases is very large, the procedure may result in considerable false positives (type I error) (Candia et al. 2008). To avoid this limitation, we also calculated the total number of anomalies detected (see Fig. 6a).

### 4.1 Anomaly detection for calling frequency

Unusual calling patterns provide a measure of behavioral response to storm severity. In the timeline, several clusters of calling frequency anomalies were observed (Fig. 6). The first occurred on 14 April, the Bengali New Year, followed by a drop the following day. A second cluster occurred on 9 May around an infamous series of national protests, dubbed “the Siege of Dhaka” in which several dozen people were killed across the nation in a series of violent protests. The next two clusters coincided with Mahasen, which made landfall on 15 and 16 May, and a cold front, which flooded the southern coast between 30 May and 1 June. Finally a large spike on 25 June coincided with Shab-e-Barat, an important religious festival, when people commonly call their relatives.



**Fig. 6** The temporal and spatial distribution of anomalies in calling frequency (**a, b**), recharge anomalies(**c, d**), and flow anomalies (**e, f**). The threshold for detection was set at three standard deviations from the mean of baseline. In the flow anomaly analysis, the restriction for including flows in the investigation was that they were positive for each day in the timeline

The most profound and longest lasting set of anomalies coincided with Mahasen. With few exceptions, calling frequency anomalies were concentrated along the vulnerable coastline, and in areas where the storm made landfall (Fig. 6b). These anomalies spatiotemporally coincided with cyclone landfall, and indicated when and where behavioral response to the physical cyclone were strongest.

Typically, post-cyclone damage assessments, which determine the form that disaster relief should take, rely on the reporting of damage by local officials. These rapid reports typically form the basis for selection of areas in which more detailed information on impacts and needs may be collected through household surveys. However systematic biases and delays can be introduced at various stages of the assessment process (Hallegatte and Przyluski 2010) due to limited capacity of responding agencies in many resource constrained settings. This analysis indicates how mobile network data could be used to overcome potential biases in the site selection portion of post-cyclone damage assessments by indicating when and where impacts have occurred.

## 4.2 Anomaly detection relating to recharge behavior

Mobile recharge anomalies differ from other anomalous behaviors detected in the timeline in two ways: they are almost exclusively concentrated around the cyclone impact zone, and occurred before the event, indicating foreknowledge of the coming cyclone and preparation for its impacts. As with calling frequency, recharge anomalies were concentrated along the vulnerable coastline, where the cyclone first made landfall (Fig. 6). But whereas calling frequency anomalies were detected widely across the region, mobile recharge anomalies were concentrated in Kalapara and Patharghata (Fig. 6d), areas noted as pockets of exceptional vulnerability within this already vulnerable landscape (Ahamed et al. 2012), where impacts of Cyclones Sidr (2007) and Aila (2009) were most severe. This suggests that anomalous preparation behaviors can reveal the areas where people perceive themselves to be vulnerable (Fig. 6d).

Alternative explanations include that these were areas where people had greater access to recharge vendors and spending money, however if this were the case, similar anomalies would have been observed in other urban centers of similar size. It is plausible that people undertook other anticipatory actions in areas where recharge anomalies were detected, but further study is required to link mobile recharges to overall disposition toward cyclone preparedness.

## 4.3 Anomaly detection relating to population movements

Forms of human mobility that deviate from normal patterns can also be detected from mobile network data (Fig. 6). To identify anomalous flows, we investigated daily flow between each pair of locations (unions) during the whole period. If the daily flow between unions A and B exceeded three standard deviations from the mean for that weekday during the study period, a signal was generated for both unions A and B. Since a substantial number of location pairs normally have low flows, false alarms could result from small absolute increases. While anomalous decreases in daily flow were less likely to produce false alarms, there is potentially more noise in increases due to the small flows which normally pass between some locations. For simplicity, to decrease noise and to obtain a conservative measure of flow anomalies, we therefore focused on pairs of locations having non-zero flow during all days during the whole period.

As with call frequency anomalies, the flow anomalies detected in this analysis corresponded with the significant events in the time line: Bengali New Year (14 April), nation-wide protests (8 May), Cyclone Mahasen (16 May), and heavy rainstorms (30 May) (Fig. 6e). The largest cluster of anomalous flow increases coincided with the Bengali New Year, where very few anomalous flow decreases were simultaneously observed. Both anomalous



flow increases and decreases were apparent between 15 and 19 May, before and after Mahasen struck, nevertheless the frequency of anomalous increases was eight times that of decreases. Note that the areas with most anomalous flow events coincided very well with the area in which rainfall intensity was the highest during cyclone landfall (Fig. 2).

## 5 Conclusion

In this paper, we show how data from mobile networks provides insights into behavioral responses to Cyclone Mahasen and its impacts. We show that anomalous patterns of calling frequency are correlated with rainfall intensity at local scales, likely providing a defined spatiotemporal measure of users' physical exposure to the storm. We show that mobile recharge purchases increase in vulnerable impact zones before landfall, representing preparations for potential environmental hazards. We also identify anomalous patterns of mobility during evacuation and storm landfall, indicating how people respond to storm forecasts and early warnings. The analysis is in agreement with the official joint needs assessment, which saw little evidence of mass displacement. We also show how, in future applications, anomalous flows of SIM cards between mobile towers can provide a much needed audit of the effectiveness of forecasting and early warnings systems, and indicate the new locations of displaced people. Rapid, cost-effective and accurate tools for assessing the effectiveness of early warning systems, and indicating the location of displaced people are currently in short supply.

Based on comparisons with rainfall measurements at landfall, and considering the substantial weakening of cyclones as they pass over land, calling frequency and population movement anomalies seemed to provide the best proxy indicators for cyclone impacts among those evaluated. Traditional methods for assessing cyclone impacts and human behavioral responses have well known limitations (Hallegatte and Przulski 2010), and the anomaly detection technique applied to mobile network data presented here (building on work of Blumenstock et al. 2011; Candia et al. 2008; Dobra et al. 2014; Kapoor et al. 2010; Pawling et al. 2007; Sundsøy et al. 2012, and Young et al. 2014), may overcome some of these challenges, and demonstrates the potential value of mobile network data as a complement to current cyclone impact assessment tools. Specifically, the spatiotemporal distributions of anomalous usage activity could be used to improve the timeliness and cost-effectiveness of cyclone impact assessments. Data from mobile networks may be very useful as a tool to prioritize locations in which rapid needs assessments are performed after cyclone landfall, with the potential to drastically reduce the time to reach those most in need.

While the study provided a robust analysis of the behavior of Grameenphone subscribers, the primary limitations of the study involved the representativeness of data for the general population. However, the general features of behavior change that we found to be most useful, i.e. sharp increases in calling frequency and changes in mobility, may well result independently of mobile operator and are likely to reflect natural human responses to shocks. Likewise, the study concentrated on Mahasen, a relatively Cyclone, which despite maximum rainfall of 68 mm/h dissipated quickly. Findings cannot be generalized about larger, more energetic cyclones, where storm surges and flooding can cause greater destruction. Finally, other causes of increased calling frequency and mobility than those indicating a need for post-disaster assistance may exist after a disaster, and thus network data should, at our present level of understanding, be used as a complement to, not a replacement for, other information sources.

To overcome these limitations and to better understand the effects of multiple types of environmental disruption, future work should incorporate mobile network data covering longer time spans. Longitudinal household measures of storm impacts and improved environmental impact models can provide external validation of the spatiotemporal patterns of anomalous usage that are apparent in the mobile network data. Additionally, as we illustrate in the Supporting Information (S2), analysis of other aspects of network function, such as service interruptions, which do not convey information on human behavior, may still provide a proxy for spatiotemporal damage to infrastructure.

Detecting anomalous usage patterns from mobile network data is a promising avenue for researching human behavioral responses to impacts associated with climate change across large spatiotemporal scales. Data from mobile networks may become an important tool for prioritizing areas for rapid needs assessments following cyclones.

**Acknowledgments** Grameenphone provided data, analysis and dedicated support to the project. The Bangladesh Telecommunications Regulatory Commission provided guidance and assured that the project adhered to regulatory standards. Rockefeller Foundation financed the work. DW and MN acknowledge Munich Re Foundation, United Nations University-Institute for Environment and Human Security, and the International Centre for Climate Change and Development for salary support. XL acknowledges the Natural Science Foundation of China under Grant Nos. 71301165 and 71522014. AJT is supported by Wellcome Trust Sustaining Health Grant (106866/Z/15/Z) and funding from the Bill and Melinda Gates Foundation (OPP1106427, 1032350, OPP1134076, OPP1117016). LB was funded by the Swedish Research Council. Stephen Roddick helped to assemble the project team.

**Open Access** This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

## References

- Adger WN et al. (2014) Human security Climate Change:755–791
- Ahamed S, Rahman MM, Faisal MA (2012) Reducing cyclone impacts in the coastal areas of Bangladesh: a case study of Kalapara Upazila. *J Bangladesh Inst Plan* 5:185–197 ISSN 2075:9363
- Associated Press (2013) Cyclone Mahasen hits coast of Bangladesh. Telegraph, London
- Auffhammer M, Hsiang SM, Schlenker W, Sobel A (2013) Using weather data and climate model output in economic analyses of climate change Review of Environmental Economics and Policy:ret016
- Bagrow JP, Wang D, Barabasi A-L (2011) Collective response of human populations to large-scale emergencies. *PLoS One* 6:e17680
- Bengtsson L, Lu X, Thorson A, Garfield R, von Schreeb J (2011) Improved response to disasters and outbreaks by tracking population movements with mobile phone network data: a post-earthquake geospatial study in Haiti. *PLoS Med* 8. doi:10.1371/journal.pmed.1001083
- Bengtsson L et al. (2015) Using mobile phone data to predict the spatial spread of cholera. *Sci Report* 5. doi:10.1038/srep08923
- Bindoff NL et al. (2013) Detection and attribution of climate change: from global to regional
- Black R, Bennett SR, Thomas SM, Beddington JR (2011) Climate Change: Migration as adaptation. *Nature* 478: 447–449
- Black R, Kniveton D, Schmidt-Verkerk K (2013) Migration and climate change: Toward an integrated assessment of sensitivity. In: *Disentangling Migration and Climate Change*. Springer, pp 29–53
- Blumenstock JE, Fafchamps M, Eagle N (2011) Risk and reciprocity over the mobile phone network: evidence from Rwanda Available at SSRN 1958042
- Brouwer R, Akter S, Brander L, Haque E (2007) Socioeconomic vulnerability and adaptation to environmental risk: a case study of climate change and flooding in Bangladesh. *Risk Anal* 27:313–326



- Candia J, González MC, Wang P, Schoenharl T, Madey G, Barabási A-L (2008) Uncovering individual and collective human dynamics from mobile phone records. *J Phys A Math Theor* 41:224015
- Dobra A, Williams NE, Eagle N (2014) Spatiotemporal detection of unusual human population behavior using mobile phone data arXiv eprint archive:1411.6179–1411.6179
- Field CB et al. (2014) Summary for policymakers Climate change 2014: impacts, adaptation, and vulnerability Part a: global and sectoral aspects Contribution of working group II to the fifth assessment report of the intergovernmental panel on climate change:1–32
- GSMA Intelligence (2015) The mobile economy. GSMA, London
- Gutro R, Pierce H (2013) NASA sees cyclone Mahasen hit Bangladesh. Goddard Space Flight Center, Greenbelt
- Hallegatte S, Przyluski V (2010) The economics of natural disasters: concepts and methods World Bank Policy Research Working Paper Series, Vol
- Hurrell JW (1995) Decadal trends in the North Atlantic Oscillation: regional temperatures and precipitation. *Science* 269:676–679
- Jones PD, New M, Parker DE, Martin S, Rigor IG (1999) Surface air temperature and its changes over the past 150 years. *Reviews of Geophysics-Richmond Virginia then Washington* 37:173–200
- Kapoor A, Eagle N, Horvitz E (2010) People, Quakes, and Communications: Inferences from Call Dynamics about a Seismic Event and its Influences on a Population. In: AAAI Spring Symposium: Artificial Intelligence for Development
- Karatepe IA, Zeydan E (2014) Anomaly Detection In Cellular Network Data Using Big Data Analytics. In: European Wireless 2014; 20th European Wireless Conference; Proceedings of, 2014. VDE, pp 1–5
- Knutson TR et al. (2010) Tropical cyclones and climate change. *Nat Geosci* 3:157–163
- Lenoir J, Gégout J, Marquet P, De Ruffray P, Brisse H (2008) A significant upward shift in plant species optimum elevation during the twentieth century. *Science* 320:1768–1771
- Lu X, Bengtsson L, Holme P (2012) Predictability of population displacement after the 2010 Haiti earthquake. *Proc Natl Acad Sci U S A* 109:11576–11581. doi:[10.1073/pnas.1203882109](https://doi.org/10.1073/pnas.1203882109)
- Lucht W et al. (2002) Climatic control of the high-latitude vegetation greening trend and Pinatubo effect. *Science* 296:1687–1689
- Lucini B, Hatt T (2014) GSMA Intelligence Report: Country Overview: Bangladesh. London
- Mann ME, Bradley RS, Hughes MK (1998) Global-scale temperature patterns and climate forcing over the past six centuries. *Nature* 392:779–787
- Martin M, Billah M, Siddiqui T, Abrar C, Black R, Kniveton D (2014) Climate-related migration in rural Bangladesh: a behavioural model. *Popul Environ* 36:85–110
- McGranahan G, Balk D, Anderson B (2007) The rising tide: assessing the risks of climate change and human settlements in low elevation coastal zones. *Environ Urban* 19:17–37
- Nadiruzzaman M (2013) Personal communication with Secretary Mesbah Alam of Bangladesh's Ministry of Disaster Planning and Response
- Olsson L et al. (2014) Livelihoods and poverty Climate Change:793–832
- Palmer PI, Smith MJ (2014) Earth systems: Model human adaptation to climate change. *Nature* 512:365–366
- Pastor-Escuredo D et al. (2014) Flooding through the lens of mobile phone activity. In: Global Humanitarian Technology Conference (GHTC), 2014 IEEE, 10–13 Oct. 2014, pp 279–286. doi:[10.1109/ghtc.2014.6970293](https://doi.org/10.1109/ghtc.2014.6970293)
- Pawling A, Chawla NV, Madey G (2007) Anomaly detection in a mobile communication network. *Comput Math Organ Theory* 13:407–422
- Penning-Rowsell EC, Sultana P, Thompson PM (2013) The 'last resort'? Population movement in response to climate-related hazards in Bangladesh. *Environ Sci Pol* 27:S44–S59
- REACH Initiative (2013) Tropical storm Mahasen final assessment report (Phase 3 Shelter, WASH, Early Recovery Detailed Assessment) Bangladesh 2013. Dhaka, Bangladesh
- Samarajiva R (2005) Policy commentary mobilizing information and communications technologies for effective disaster warning: lessons from the 2004 tsunami new media. *Society* 7:731–747
- Seneviratne SI et al. (2012) Changes in climate extremes and their impacts on the natural physical environment Managing the risks of extreme events and disasters to advance climate change adaptation:109–230
- Serreze M et al. (2000) Observational evidence of recent change in the northern high-latitude environment. *Clim Chang* 46:159–207
- Stenseth NC, Mysterud A, Ottersen G, Hurrell JW, Chan K-S, Lima M (2002) Ecological effects of climate fluctuations. *Science* 297:1292–1296
- Stroeve J, Holland MM, Meier W, Scambos T, Serreze M (2007) Arctic sea ice decline: Faster than forecast *Geophysical research letters* 34
- Sundsøy PR, Bjelland J, Canright G, Engø-Monsen K, Ling R (2012) The activation of core social networks in the wake of the 22 July Oslo bombing. In: Advances in Social Networks Analysis and Mining (ASONAM), IEEE/ACM International Conference on, 2012. IEEE, pp 586–590

## Supporting Information:

### Detecting climate adaptation with mobile network data in Bangladesh: Anomalies in communication, mobility and consumption patterns during Cyclone Mahasen

Xin Lu<sup>a,b,c</sup>, David J. Wrathall<sup>d</sup>, Pål Roe Sundsøy<sup>e</sup>, Md. Nadiruzzaman<sup>f,g</sup>, Erik Wetter<sup>b,h</sup>, Asif Iqbal<sup>e</sup>, Taimur Qureshi<sup>e</sup>, Andrew Tatem<sup>b,i</sup>, Geoffrey S. Canright<sup>e</sup>, Kenth Engø-Monsen<sup>e</sup>, Linus Bengtsson<sup>a,b\*</sup>

<sup>a</sup>Department of Public Health Sciences, Karolinska Institutet, Stockholm, Sweden;

<sup>b</sup>Flowminder Foundation, Stockholm, Sweden;

<sup>c</sup>College of Information System and Management, National University of Defense Technology, Changsha, China;

<sup>d</sup>Oregon State University, College of Earth, Ocean and Atmospheric Sciences, Corvallis, Oregon, USA;

<sup>e</sup>Telenor Research, Oslo, Norway;

<sup>f</sup>Department of Geography, University of Exeter, Exeter, UK;

<sup>g</sup>International Centre for Climate Change and Development, Dhaka, Bangladesh;

<sup>h</sup>Stockholm School of Economics, Stockholm, Sweden;

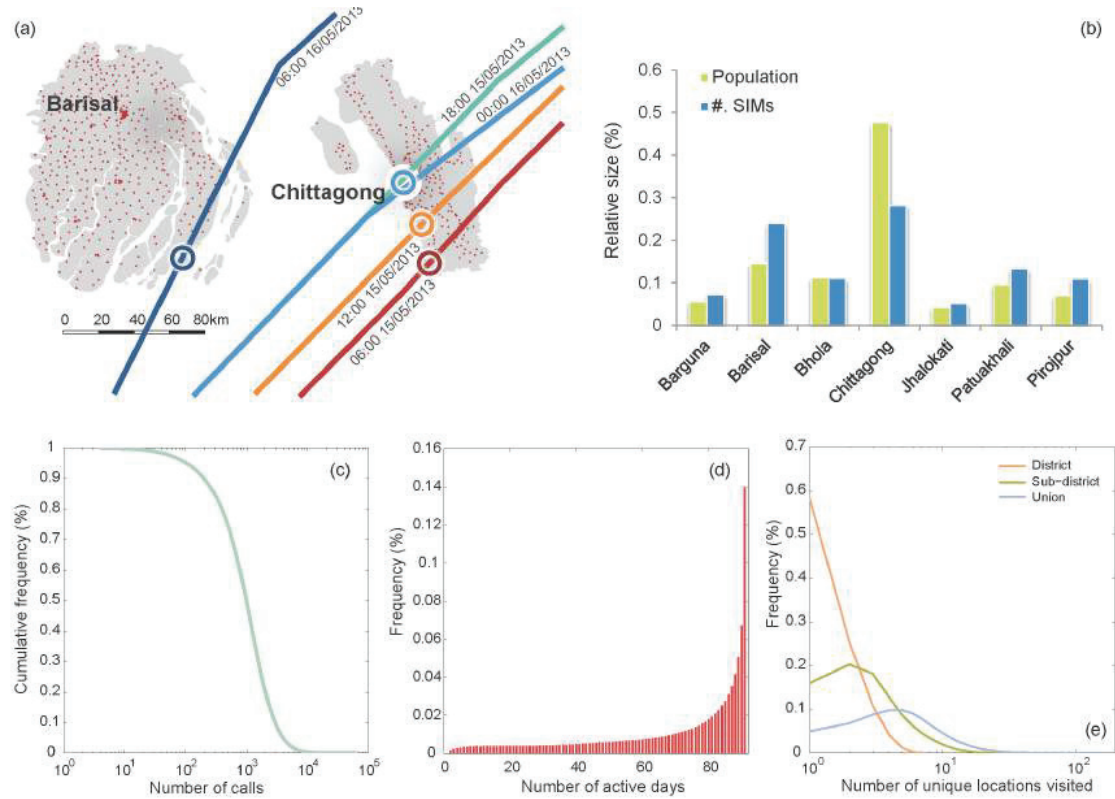
<sup>i</sup>WorldPop, Department of Geography and Environment, University of Southampton, Southampton, UK.

\*Linus Bengtsson: [linus.bengtsson@flowminder.org](mailto:linus.bengtsson@flowminder.org)

#### S1. Dataset Characteristics

De-identified call detail records (CDRs) from 5.1 million Grameenphone users were collected between 1 April and 30 June of 2013 in the Barisal Division and Chittagong District of Bangladesh. The data begins six weeks before the landfall of Cyclone Mahasen (16 May 2013) and continues for six weeks after landfall (1 April to 30 June 2013) (Fig. S1a). CDRs are compiled by network operators principally for the purposes of billing customers for their use of the network. De-identified data entries include information on the time of the call, the mobile phone tower used and the duration of call, and can thus be used to indicate the geographical position and movements of users. To limit potential biases resulting from subscriber churn, and new users entering the dataset due to impacts of the storm, we limited the study to SIM cards that had placed at least one call before the cyclone landfall (16 May); and also made at least one call in the last ten days of the data collection period (21-30 June).

A distribution of 986 towers across the Barisal Division and Chittagong District forms the basis for our spatiotemporal analysis of calling frequency, mobility and top-up behaviors. We assigned users a position within the network based on the tower through which his or her most recent call was routed (Fig. S1a). The coverage area of a mobile tower (BTS, base transceiver station) can range from a few hundred meters in a city, up to tens of kilometers in rural areas. The location of tower positions were further scrambled within 200 meters. The average and maximum distance of each tower and its nearest neighbor tower, were 2.1 km and 16.8 km, respectively.



**Fig. S1** Study area, distribution of users, distribution of calls, number of active days, and unique locations visited. a) Map of study area, including tower map and cyclone path. Colored lines represent the forecasted versus actual cyclone paths. Forecasts predicted landfall over Chittagong, but on its final approach, Mahasen passed over Barisal. Actual versus predicted storm paths allow us to distinguish anticipatory and actual behaviors. b) Distribution of population and number of SIM cards at district level. Population data are taken from the 2011 census, and SIM cards' locations are determined by the district in which they appeared most frequently during evening (10:00pm-6:00am) during the study period. c) A distribution of the number of calls per user shows that 90% of users made at least 200 calls during the three months period, and 50% of users made more than 1000 calls during the period. d) The number of active days on the network shows that 89% of users were active on 30 days, and 50% of users were active on at least 80 days. e) A distribution of unique locations visited shows that 95% of users visited more than one union.


A comparison of number of SIM cards to population shows high agreement between the number of subscribers and the census figures (Fig. S1b). Deviations are seen in Chittagong district. In Bangladesh, household possession of a mobile phone grew from 78% of all households in 2011 to 89% in 2014, exceeding usage in other countries of similar socioeconomic profiles (National Institute of Population Research and Training (NIPORT) 2015). Users from the Grameenphone network, the largest in Bangladesh, numbered at 42 million in 2013, and comprised 61% of all mobile users nationwide (Telenor 2013). In this dataset, 90% of users made more than 200 calls during the study period and were active 80 days or more out of the 90 day study period (Fig. S1c, d). Likewise, SIM mobility is high: 95% of users appeared in more than one union, 84% in more than one sub-district, and 42% in more than one district (Fig. S1e). For the purposes of this study, we assumed that behaviors exhibited in CDRs are representative for the entire population. While earlier research support this notion, more studies are needed to adequately understand the representativeness of mobile phone users in Bangladesh.

## S2. Quantifying impacts and infrastructure resilience

The functioning of mobile network towers potentially provides an additional proxy for cyclone damage. We undertook an analysis of network function during the study period, investigating towers that were inactive during Mahasen. Because of the potential for damaged equipment and interruption of services, mobile operators invest heavily in network resilience. Therefore in the most powerful cyclones, disabled towers could indicate where damages are concentrated, and could provide a proxy for other infrastructural damages sustained.

In the analysis of towers with no calling activity, we see the same four events that appear in other sections. Focusing on Mahasen, we evaluated the number of inactive towers in Mahasen's impact zone, i.e., towers that registered zero calls, calculated on an hourly basis (Fig. S2). The Grameenphone network held up very well. During Cyclone Mahasen, only 60 towers went offline during landfall followed by 120 towers going offline during the course of the following day. It is likely that rather than suffering damages, towers went offline the following day as power sources failed and reserve batteries became depleted. Notably, patterns of "offline" towers between Mahasen and a heavy downpour that occurred between 30 May and 2 June were remarkably similar (Fig. S2).

In contrast to anomalies in calling frequency, recharges and mobility, tower function does not represent behavioral information, but may provide a potential indicator on the distribution of environmental impact to infrastructure. Mahasen was a relatively weak storm, and as such registered only minor disruption, however during major events, data showing inactive towers or loss of tower function may also indicate the locations where damages are concentrated, helping focus impact assessment on the most damaged sites. Further research is required to make meaningful inferences about the relationship between network function and infrastructural damage.



**Fig. S2** Number of towers with zero calls over the study period. For all filtered SIM cards, the number of calls received by each tower in each hour is calculated. Towers are classified as inactive during any particular hour if they received zero calls. There are four major periods when the number of inactive towers increased: 1) the Bengali New Year, 2) a national demonstration, 3) Cyclone Mahasen, and 4) a severe rainstorm.

**References:**

- National Institute of Population Research and Training (NIPORT), Mithra and Associates, and ICF International (2015) Bangladesh Demographic and Health Survey 2014: Key Indicators. NIPORT: Dhaka, Bangladesh and Rockville, Maryland.
- Telenor, (2013) Grameenphone, Bangladesh: In Bangladesh, Grameenphone is the leading provider of mobile communications services. <http://www.telenor.com/investors/company-facts/business-description/grameenphone-bangladesh/> Accessed 9 June 2015.

## **Paper 6**

---





## **Paper 7**

---

