Enabling the responsible use of personal data with targeted differential privacy

Joshua Blumenstock and Nitin Kohli

Mobile Phones and Mobile Money

Mobile phones are now widespread in LMICs

- 5.4 Billion mobile subscribers (68% penetration; 43% SSA, 70% LAM, 62% AP) globally
- 4.4 Billion mobile internet users

So too are services built on the mobile network

- Mobile money: 763m registered accounts in sub-Saharan Africa
- Loans, payments, tele-medicine, insurance, pay-as-you-go financing, agricultural extension…

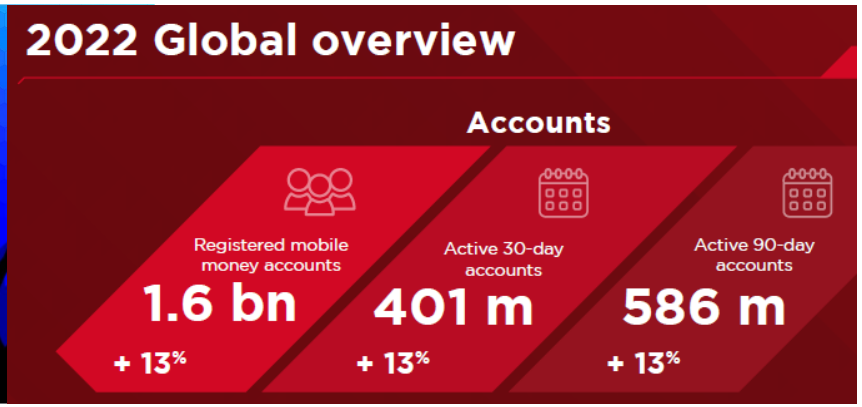

Personal Data from Emerging Markets

These devices and services are generating vast troves of data on personal behavior

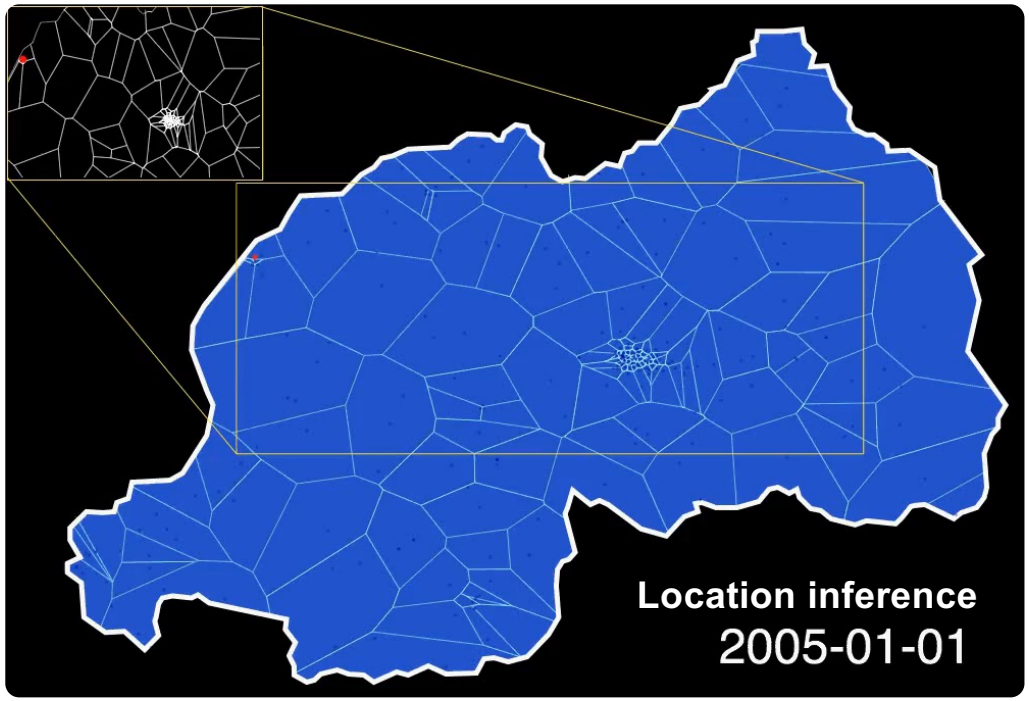

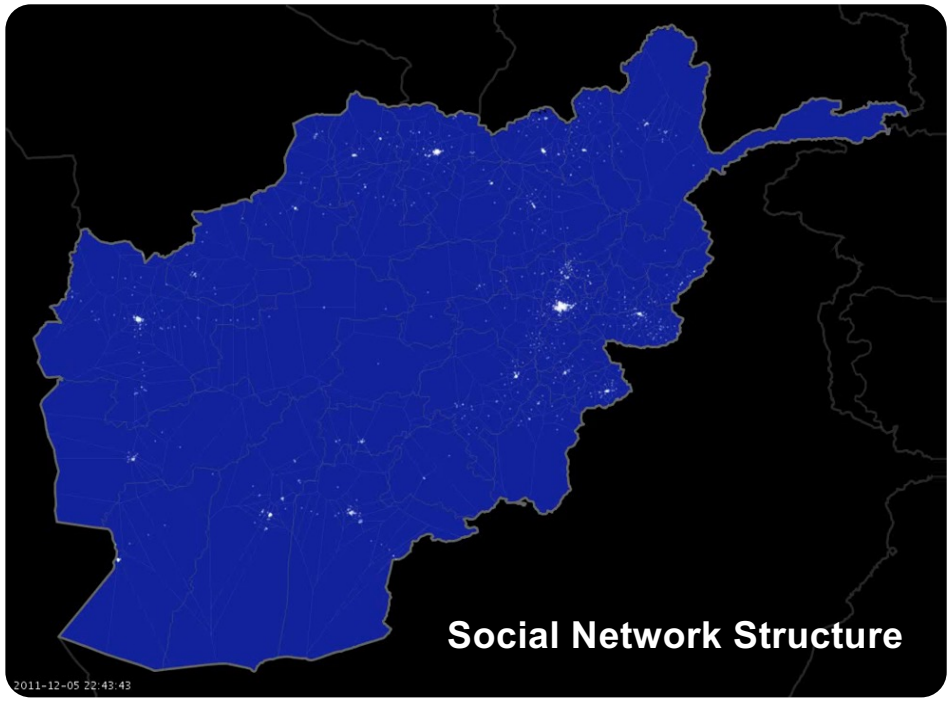

Applications of Personal Data

The data, in turn, are enabling rapid innovation. Two examples:

Digital Credit / Instant Loans

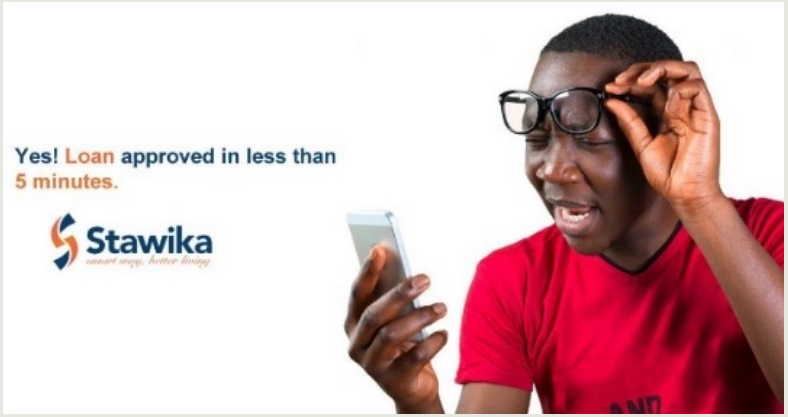

- **Science:** Machine learning and mobile phone data can predict *who will repay a loan*
- **Product:** This makes it possible to lend to people without formal financial histories
- **Impact:** 27% of Kenyans have a digital credit loan

Targeting of Humanitarian Aid

- **Science:** Machine learning and mobile phone data can predict *who is living in extreme poverty*
- **Product:** This makes it possible to transfer money to people during a humanitarian crisis
- **Impact:** 150,000 extreme poor received *Novissi* cash payments

Concerns with personal data privacy

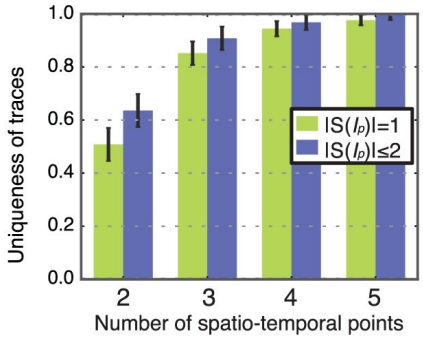

But, the data can reveal sensitive information

- Age, gender, race, sexual orientation, location, occupation, and much more

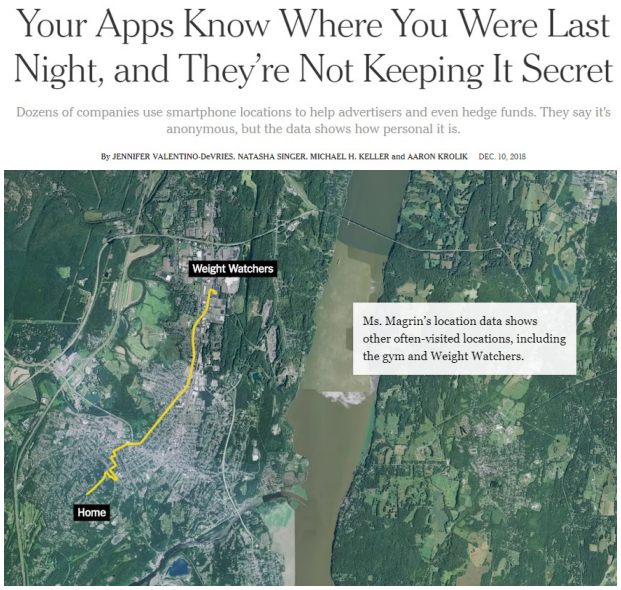

Our work: *Provably* Private Targeting

We are working on two fronts:

1. Develop algorithms that provide strict (formal) data privacy guarantees, while also facilitating downstream applications of machine learning

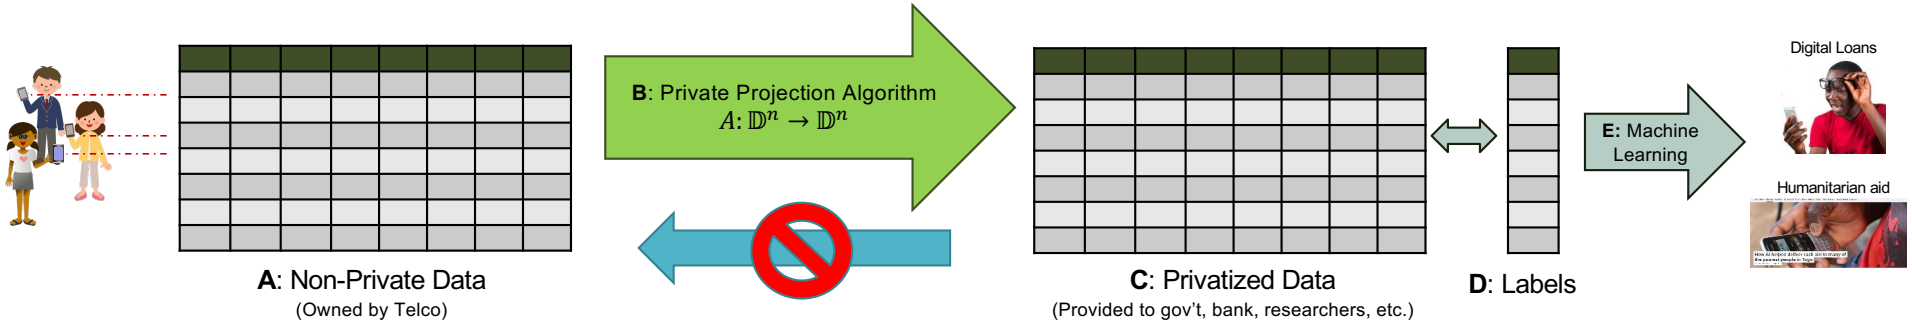

2. Characterizing the real-world tradeoffs induced by increasing privacy

Contribution 1: Targeted differential privacy

Our algorithm for *targeted differential privacy* (show proof) substantially outperforms differential privacy

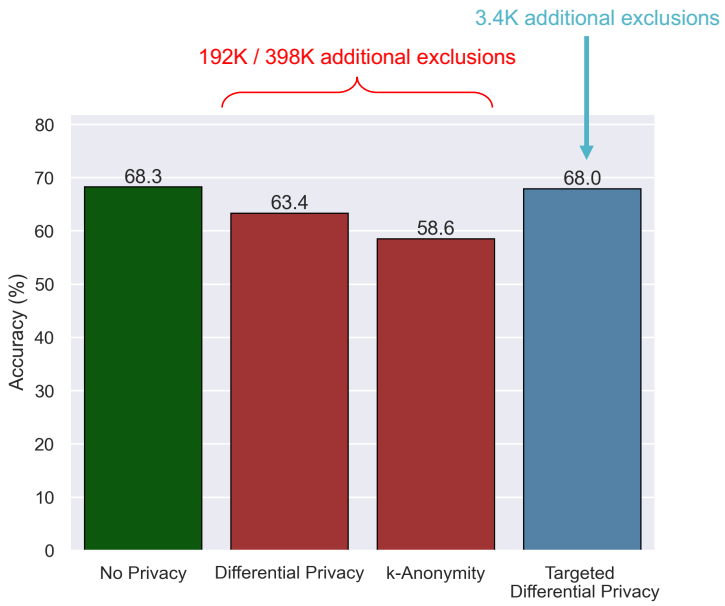

A: Targeting of humanitarian aid (Togo)

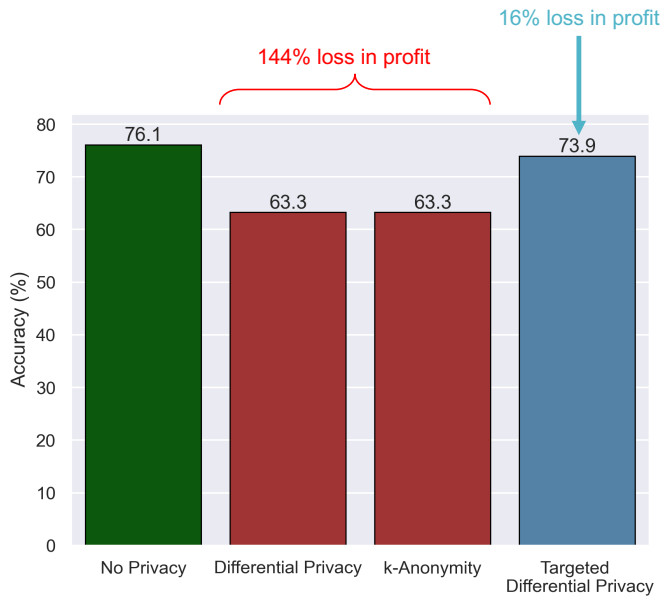

B: Digital credit loan applications (Nigeria)

Contribution 2: Tradeoffs induced by privacy

Example result: Singling out attacks in Togo's *Novissi* program

Concern: Can an adversary hack the private data to re-identify least one individual (based on unique combinations of features)?

Key idea: Privacy enhancing technologies can protect against such attacks, but their use reduces the performance: more non-poor get benefits (inclusion errors) and more poor don't get benefits (exclusion errors)

Results:

- **Relative to baseline:** Singling out protection increases by 8,480%, at the cost of 3,400 exclusion errors
- **Relative to differential privacy:** Singling out protection decreases by 12.6%; 187,500 fewer exclusion errors

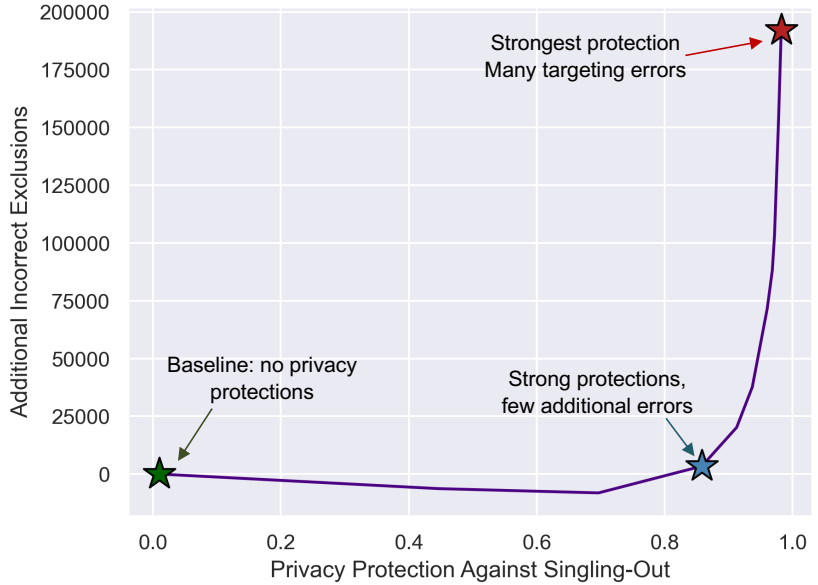

Summary

Big data is enabling powerful new applications in LMICs, but creates risks to personal privacy

- Existing privacy enhancing technologies (like differential privacy) are not well-suited to these applications
- Targeted differential privacy gives decision-makers granular control over the level of privacy

Increased privacy comes at a cost

- Our work characterizes these tradeoffs in two real-world programs
- More broadly: tries to make decisions around data privacy more legible and actionable to decisionmakers

Necessary condition for accurate targeting

Theorem: Suppose $A: \mathbb{D}^n \rightarrow \mathbb{O}$ satisfies (B, ϵ, δ)-TDP. Let $N: \mathbb{D}^n \rightarrow \mathbb{O}$ be a deterministic function with bounded range. If $\left|\left|A(X) - N(X)\right|\right|_\infty \leq \alpha$ with probability γ for all $X \in \mathbb{D}^n$, then

$$B < \frac{diam(\mathbb{D})}{\lceil \epsilon^{-1} \ln Q_w \rceil - 1}$$

where w is the 2α-packing number of $Range(N)$ under $||\cdot||_\infty$ and $Q_w = \frac{\delta + \gamma(e^\epsilon - 1)}{\delta + w^{-1}(e^\epsilon - 1)}$

Implication: $B \ll diam(\mathbb{D})$ for many practical applications in our setting, yielding an impossibility result