Research Article

SmartBrowse: Design and Evaluation of a Price Transparency Tool for Mobile Web Use

Nithya Sambasivan
Paul Lee
Greg Hecht
Paul M. Aoki
Maria-Ines Carrera
Jenny Chen
Michael Youssefmir

Google, Inc.

David Cohn

Natembea Foundation

Pete Kruskall

Kensho

Everett Wetchler

Bayes Impact

Astrid Twenebowa Larssen

Halogen AS

Abstract

Mobile data use is on the rise globally. In emerging regions, mobile data is particularly expensive and suffers from a lack of price and data usage transparency, which is needed to make informed decisions about Internet use. To measure and address this problem, we designed SmartBrowse, an Internet proxy system that (1) shows mobile data usage information and (2) provides controls to avoid overspending. In this article, we discuss the results of a 10-week study with SmartBrowse, involving 299 participants in Ghana. Half the users were given SmartBrowse, and the other half were given a regular Internet experience on the same mobile phone platform. Our findings suggest that, compared with the control group, using SmartBrowse led to (1) a significant reduction in Internet credit spend and (2) increased online activity among SmartBrowse users, while (3) providing the same or better mobile Internet user experience. Additionally, SmartBrowse users who were prior mobile data non-users increased their web page views while spending less money than control users. Our discussion contributes to the understanding of how ICTD research with emerging technologies can empower mobile data users, in this case, through increased price and usage transparency.

To cite this article: Sambasivan, N., Lee, P., Hecht, G., Aoki, P. M., Carrera, M.-I., Youssefmir, M., et al. (2015). SmartBrowse: Design and evaluation of a price transparency tool for mobile web use [ICTD2013 Special Issue]. *Information Technologies & International Development, 11*(1), 21–40.

Introduction

While the individual benefit of mobile ICT *use* in emerging regions is an active topic of research, understanding mobile ICT *spending* is equally important. For example, a nationally representative consumer survey in 17 countries in sub-Saharan Africa (SSA) showed that mobile phone spending constituted 10–26% of individual income in the lower 75% income bracket (Chabossou, Stork, Stork, & Zahonogo, 2009).

Because of the significant individual financial impact of mobile ICT use, we believe that enabling mobile users to make the best use of their ICT spending should be a key area for ICTD (information and communication technologies for development) research—and in particular, we believe that mobile *data* use should be a key area for ICTD research on emerging technologies. While many ICTD practitioners rightly focus on maximizing access for the lowest socioeconomic groups by using lowest-common-denominator mobile services (such as voice calls, SMS, and USSD [Unstructured Supplementary Service Data]), ICTD research cannot ignore the fact that mobile data use is rising rapidly. This rise cuts across a range of socioeconomic groups because mobile data can now enable substantial cost savings (e.g., over-the-top voice and messaging services such as WhatsApp can be less expensive than the carriers' voice and messaging services) as well as new, innovative services.

Price transparency is a significant barrier to making informed choices about mobile data use. A recent McKinsey study found that greater than 20% of mobile data non-users in major African cities cited lack of pricing information and control over monthly expenses as key factors for their Internet non-use (Sun Basorun, Moodley, Moraje, & Nielsen, 2012). Consumers can easily manage carrier voice and messaging costs because the units of cost and billing (minutes and messages) are familiar and clear. By contrast, the cost of Internet browsing is opaque. What is the cost to load a "web page"? What does a balance of 3 MB mean in terms of my online activities? Which websites consume the most credit? While there have been attempts to normalize per-page costs via content modification (Chava, Ennaji, Chen, & Subramanian, 2012), this creates a technically brittle, second-class browsing experience. Is there a way to enable mobile users to understand their data spending without altering content?

In this article, we present the results of a 10-week study of 299 mobile Internet users in urban Ghana. We provided participants with Android phones that enabled a full, standards-compliant Internet browsing experience and gave half of them access to an Internet proxy system called SmartBrowse. Through a variety of features, the system informed users of the cost of accessing a given web page in context; that is, *prior to and immediately after incurring that cost,* providing significant evidence that price transparency for Internet browsing can be increased without rewriting content.

The contribution of this article is twofold. First, we describe the SmartBrowse system and its effectiveness in reducing mobile data spending without negatively affecting web browsing behavior. To the best of our knowledge, our research is the first study of users in emerging regions that seeks to understand the effect of mobile data price and usage transparency *before* and *while* using the Internet. Second, we provide a detailed description of the logistics involved in running a user study that requires this degree of "new" infrastructure—studying interventions not currently pervasive or affordable to a majority of socioeconomic groups, but that are steadily increasing—deployed longitudinally with randomly selected users. Our research approach contrasts with other ICTD interventions of emerging technologies that focus on technologically enhancing an institutional or commercial worker's efficiency, e.g., NGO health workers (Ramachandran, Canny, Das, & Cutrell, 2010) or clinicians (Anokwa, Ribeka, Parikh, Borriello, & Were, 2012).

Our article is organized as follows. We highlight our baseline research findings on mobile data attitudes, which informed our software design on pricing information and protection. We describe the SmartBrowse system in detail, followed by a discussion of the research methods we used to measure the effects of the intervention. We then describe our findings from the trial, including Internet use and credit spend behavior, attitudes, and perceptions; overall satisfaction; and the feeling of control over mobile data credit. We follow up the findings with a discussion of price transparency, including empowering data users, helping mobile data non-users learn about Internet use, conducting research with urban users, and running a forward-looking study in ICTD. We conclude the article by describing our logistics and experience of setting up the SmartBrowse trial.

Background

The Cost of Mobile Service

The cost sensitivity of mobile phone and Internet users in emerging regions is a common ICTD concern. For example, researchers frequently propose low-cost access services based on "alternative" connectivity models with reduced interactivity (e.g., Isaacman & Martonosi, 2011; Pentland, Fletcher, & Hasson, 2004). However, for a growing segment of the population, the relevant question is no longer how to obtain access. Geographic access (coverage) and economic access (affordability) of commercial mobile ICTs have improved greatly over the last decade (Williams, Mayer, & Minges, 2011). Instead, the question is how to work access into routines in both opportunistic (Sambasivan & Smyth, 2010) and planned (Wyche, Smyth, Chetty, Aoki, & Grinter, 2010) ways over the course of a day or how to negotiate access from those around them (Sambasivan & Cutrell, 2012). A 2012 report on "base of the pyramid" (BoP) mobile use in Kenya indicated that even in the BoP demographic, 25% of participants reported using mobile data services (Crandall et al., 2012).

Buying Mobile Data in Emerging Regions

It is instructive to look at what is known about how people buy mobile services, focusing on SSA where related work is available.

Telecom and Development Indicators

While country-level statistics tell us relatively little about individual purchasing behavior, they do give us an important comparative context for those behaviors. This context comes from both supply-side and demand-side sources.

Supply-side (operator-reported) data. Supply-side data includes characterizations of the service plans offered in various markets, including prices (e.g., Communications Chambers, 2012; ITU, 2013; Otsuka, 2009); the number of service plan subscriptions (e.g., Williams et al., 2011); and the type of service plan subscriptions (e.g., ITU, 2013). This data is released by operators and collated by various methods: via national regulators and the ITU, shareholder reports, websites, etc.

For our purposes, the key takeaway from the supply-side data is that prepaid mobile service plans, sold in small increments, have increased the affordability of mobile service in emerging regions (e.g., Minges, 1999). In 2012 an estimated 87% of mobile subscriptions in emerging regions were prepaid (ITU, 2012); in 2013 this number was 95% in Africa (Gillet, Hatt, & Lucini, 2013). The logic here is essentially the same as that for small packet (sachet) marketing of physical goods (Prahalad, 2005): Availability in smaller sales units increases product affordability for customers who cannot easily obtain credit or save larger amounts of cash (Gillet et al., 2013). Less research exists on mobile data use, but the same logic would be expected to apply to mobile data (Donovan & Donner, 2010) and the limited operator data available suggests that consumers do prefer smaller units for prepaid mobile data (e.g., see Vodacom, 2012).

Demand-side (consumer survey) data. Demand-side data is generally obtained through surveys that assess what subscribers actually pay for mobile services (e.g., Chabossou et al., 2009), typically as a fraction of income as well as in absolute terms, and what services they actually use (e.g., Crandall et al., 2012; Stork, Calandro, & Gillwald, 2013).

For our purposes, the key takeaway from the demand-side data is that despite dramatic affordability (and adoption) gains over the last decade, spending on mobile services still consumes a large fraction of the disposable income for many SSA consumers. Mobile service costs in SSA represent double-digit percentages of income for all but the top income quartiles, unlike the low-single-digit percentages typical in developed economies (Chabossou et al., 2009). Even in highly competitive telecom markets such as Kenya, surveys reveal that the poorest subscribers still often cut back on necessities (e.g., food) to access those services (Crandall et al., 2012).

Studies of ICT Consumption in Local Markets

Prepaid credit balances for mobile service (usable for voice call "airtime," SMS, etc.) are generally "topped up" by purchasing "scratch cards" or by using a mobile payment service balance from a ubiquitous network of informal traders. Naturally, both technical mechanisms and informal practices (e.g., Murphy & Priebe, 2011; Sambasivan & Cutrell, 2012) for balance-sharing are used throughout emerging regions.

Balances can be used to pay for mobile data at a pay-as-you-go (PAYG) usage-based rate. Mobile data can also be purchased at a discounted rate in "bundles," a given allowance (in megabytes) with an expiration period (typically one, seven, or 30 days). For our purposes, the key takeaway from the limited studies of mobile purchasing is that only a minority of subscribers use the discounted bundles. For example, a 2011 survey of public venue Internet users in Cape Town, South Africa found that only 37% of teens and 32% of adults had used mobile data bundles (Walton & Donner, 2012), despite the fact that bundles are typically less expensive than PAYG. Also confirmed by Vodafone in South Africa is the fact that only one third of data users bundled their Internet (Vodacom, 2012). A 2012 survey of low-income mobile phone users in Kenya found that only 1% had any knowledge of data bundles, even though a quarter used mobile data (Crandall et al., 2012).

Usage Transparency

To date, researchers have explored ways to allow users to retrospectively view and manage Internet bandwidth use (see Chetty, Banks, Brush, Donner, & Grinter, 2012, for a study on usage practices around bandwidth caps among residential Internet subscribers in South Africa and see Chetty et al., 2010, for a bandwidth management tool). More broadly, researchers have explored ways to effect behavior change around resource consumption, such as water and electricity, in the forms of information, prompts, incentives, goal setting, and social comparison (see Froehlich, Findlater, & Landay, 2010 for an overview of feedback technologies). Our research contributes to the understanding of changes in decision making when pricing and usage transparency are made available (e.g., Narayanan, Chintagunta, & Miravete, 2007) while using the Internet by showing current balance information and providing actionable controls on a mobile phone.

Baseline Research

Mobile Data Survey in Kenya

In February 2012 we collaborated with iHub Research on an exploratory survey in Nairobi, Kenya to understand user attitudes toward mobile data pricing and usage (iHub Research, 2012b). Eighty-two participants (19 mobile data non-users) were surveyed in mixed-income sites.

The findings showed that respondents who understood the relationship between the size of a web page and the associated cost reported spending less money on mobile data.

Mobile Data Survey in Ghana

As a follow-up to the Kenya survey, we conducted a more comprehensive survey on mobile data practices and attitudes with a larger sample size (798 mobile data users and 194 mobile data non-users) in Ghana in June 2012 (iHub Research, 2012a). Participants were screened from mixed-income sites in Accra and Sunyani.

A large number of mobile Internet users (48%) spent less than US$2.50 per week on mobile data, while 28% spent less than US$2.50 per week on voice calls. Mobile data users and non-users did not have an accurate understanding of mobile Internet costs and many believed they were billed by time. Only 19% of users were able to correctly identify which mobile data activities were the most expensive. Fewer users monitored mobile Internet spending, compared to voice/SMS balances. Most (75%) users kept track of their voice/SMS balances, compared to 38% who tracked their mobile Internet balance. One third of data users accidentally spent more money on mobile data than they intended.

SmartBrowse: Price Transparency for Mobile Data

Guided by the findings from the Kenya and Ghana surveys, our motivation was to improve price transparency through three methods:

1. **Increase awareness of mobile data spending among users:** Allow users to track their credit balance and learn about web page costs as they browse.
2. **Protect users from unexpected spending:** Alert users before browsing expensive websites to prevent unexpected overspending, allowing them to decide whether to continue to visit the website.

3. **Allow users to top up easily:** Provide a standalone on-screen element to easily check and top up their credit balance, avoiding the difficult-to-use balance check through USSD (short code).

System Design

Android devices were selected for ease of control over the OS and installed apps. A dual-SIM phone was chosen to allow set-up of SIMs for data (provided to participants as part of the study, solely for data access) and voice (so that users could continue to receive voice calls and messages using their existing phone number without using multiple phones). Phone software was preconfigured with customized browser settings (e.g., proxy settings) and home screen web shortcuts. Mobile device management software restricted user access to apps and directed all Internet use through the browser, with the exception of WhatsApp, which was used to collect user feedback.

Architecture

SmartBrowse's price transparency features are entirely implemented using a web proxy server run on Google Compute Engine. Every web (URL) request from the Android devices is routed through the proxy server. The proxy server estimates the expected size of the document at that URL using historical data, and this expected document size is used to calculate the cost of loading it. Based on these estimated page costs and the user's remaining mobile data balance, the proxy provides customized spending alerts (as detailed in the next subsection).

We created a simple billing system that managed the individual mobile data balances and included a web-based top-up facility. In this trial, the user was always charged for the page cost as estimated by the proxy server prior to page load, even if this estimated page cost was found to be inaccurate after the page was loaded (i.e., the current page size did not match the estimate from historical data due to dynamic changes). Since a specific goal of the trial was to evaluate the effectiveness of page cost feedback and since cost estimate inaccuracy might reduce the value of the feedback, the trial users were not penalized for any such inaccuracy. However, the actual page sizes were logged for later analysis.

Since the study's purpose was to measure the impact of price transparency on user behavior directionally, our system calculated costs based on simple overall web page sizes. Future studies could consider the complexity of web pages when drawing in various elements from different servers (Butkiewicz, Madhyastha, & Sekar, n.d.; Wang, Balasubramanian, Krishnamurthy, & Wetherall, 2013).

The estimated cost-per-byte was determined as follows. Since most respondents (85%) in our pre-study Ghana survey used PAYG for data, we used a typical PAYG rate in July 2012 (US$0.035/MB) as the base cost. Applying a compression file-size savings of 59%, as reported by a popular third-party web proxy (Opera Turbo for Opera Mobile; von Tetzchner, 2010), the net cost is roughly US$0.015/MB (in terms of the original, uncompressed size of the web page recorded in our historical data).

For this trial, we further opted not to pre-render, cache, or compress web page data (Chava et al., 2012; Chen & Subramanian, 2013; Wang, Shen, & Wetherall, 2013). These would be possible in future trials, but each mechanism can affect user perception of cost, latency, and data freshness. Since a specific goal of the trial was to understand user perceptions of the price transparency mechanism itself, these mechanisms (with their potentially confounding effects) were omitted.

Features

SmartBrowse consisted of four main visible features (shown in Figure 1 as wireframes). Our features were designed to provide ongoing cost education and protective controls. We kept them simple enough for both new and existing Internet users to understand.

- **Balance bar:** A persistent balance bar displayed the user's current data balance in local currency over all pages [Figure 1(a)], allowing them to check their balance without effort. The balance bar also provided a link to the top-up page, described below.
- **Page costs for Google.com search results:** Every Google search results page displayed an estimated cost of following the individual search result links. Page costs were displayed below the web snippets [Figure 1(b)]. Our motivation was to educate the user about the costs of various search results.

- **"SmartAlerts":** In two cases, the SmartBrowse proxy inserted interstitial web pages that were loaded when the user clicked on links. To prevent inadvertent overspending, an "expensive page" SmartAlert appeared before loading pages with an estimated cost that exceeded a preset threshold [Figure 1(c)], allowing the user to "Continue" (incur the cost) or go "Back" (not incur the cost). Additionally, the user could customize the SmartAlert threshold, which was set by default at 3 Gp.[1] The "top-up" SmartAlert appeared when the user's balance was estimated to be too low to successfully load the requested page, allowing the user to top up or go back [Figure 1(d)].
- **Top-up page:** An online top-up page, reachable through the balance bar and home screen shortcuts, allowed the user to check their balance without relying on additional mechanisms such as USSD. The top-up page allowed the user to top up by entering custom scratch card codes (details about this in the "Running a Trial" section). Using the top-up page was free.
- **Web page shortcuts:** In addition to the system features above, shortcuts to popular websites (based on an analysis of most-visited sites from Ghana) were added to the home screen. As previously mentioned, a link to the top-up page was also placed on the home screen.

Research Design

The research trial lasted 10 weeks during the Fall 2012 university term (mid-September to mid-November). We aimed to recruit 300 participants; in total, we recruited 299 participants, whose characteristics are summarized in Table 1. Nearly all participants completed the study (n = 282), although a few phones were lost or stolen (n = 7) and some participants dropped out for other reasons (n = 10).

Our priority in screening participants was to select prior users of mobile data service. As participants were recruited through university contacts, most were students with low to middle personal income levels (their overall socioeconomic status may well have been higher).

Participants were randomly assigned to control and treatment groups. The control group could use regular Android phone features and view the top-up page (which was required to continue using the phone's data SIM). In addition, the treatment group could view and use the SmartBrowse features.

Research Methods

Baseline Survey

All participants completed a baseline survey before receiving their phones. Questions focused on mobile device ownership and use; airtime and mobile data spend; and a cost awareness exercise (to test the users' ability to assess websites based on the size of the webpage).

Midtrial Checkpoints 1 and 2—Survey + Individual Interviews

In weeks 3 and 7, all participants completed additional surveys. Midtrial survey questions focused on overall and feature-level satisfaction and usefulness, understanding of SmartBrowse, top-up behavior, perceptions of management of credit, and the cost awareness exercise. Thirty participants were interviewed. Interview questions focused on phone tours and SmartBrowse features.

Trial Exit Checkpoint—Survey + Focus Groups

In week 10, all participants completed an exit survey, which focused on satisfaction, perceptions of management of credit, future phone purchases, and the cost awareness exercise. Twenty-five participants in the treatment and control groups participated in focus group discussions.

Post-Trial Follow-Up Survey

Three months after the trial's end, we conducted a follow-up survey to understand how SmartBrowse had impacted the participants, especially around new phone purchases, mobile data plans and use, and cost awareness. All 299 participants were contacted, and 126 responses (63 control, 63 treatment) were received by e-mail or voice calls.

---

1. Ghanaian currency notations, GH₵ (Ghana cedi, US$0.51 at the time) and Gp (Ghana pesewa, or GH0.01), are used in the remainder of this article.

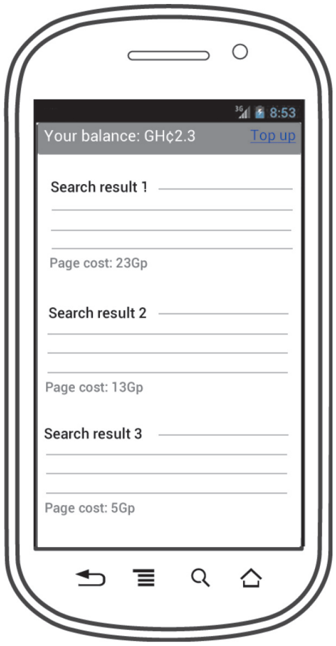

(a) Google.com page costs for search results and Balance bar

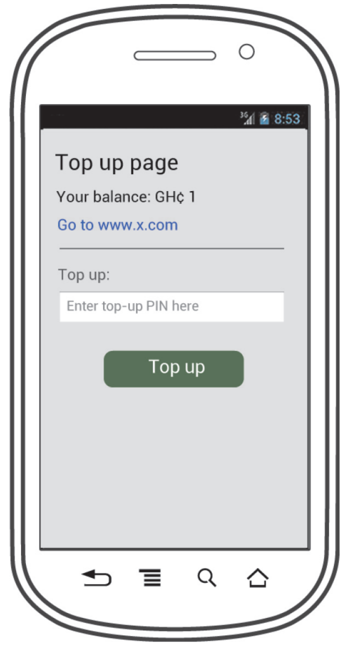

(b) Top-up page

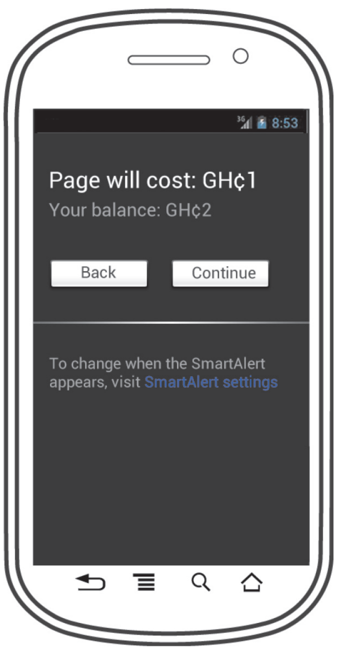

(c) "Expensive page" SmartAlert (interstitial)

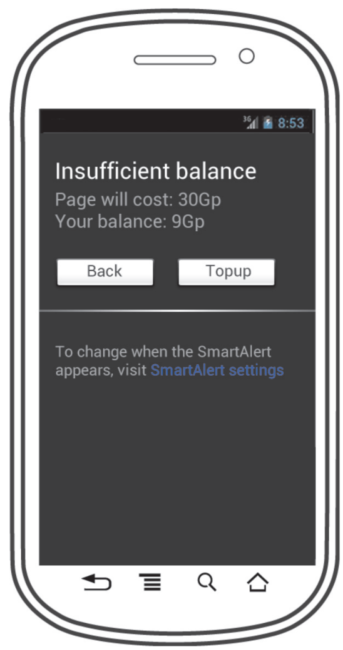

(d) "Top-up required" SmartAlert (interstitial)

Figure 1. Visible SmartBrowse features.

Table 1. Characteristics of Recruited Participants.

		Control	Treatment
Total		148	151
Prior mobile data user?	Prior users	106	117
	Prior non-users	42	34
Occupation	Students	106	121
	Nonstudents	42	30
Gender	Male	72	73
	Female	76	78
Income bracket[1]	Low	50	56
	Medium	77	79
	High	21	16

1. These descriptive brackets (low/medium/high) were set differently for the student (< GH₵200/GH₵200–400/ > GH400) and nonstudent (< GH₵1,000/GH₵1,000–3,000/> GH₵3,000) occupation groups, adjusting for differences in income levels and spending requirements between the groups. The heuristic nature of surveying individual discretionary income has been well documented in the development literature (see Haughton & Khandker, 2009; Moore, Stinson, & Welniak, 2000 for reviews).

Logs Analysis

In addition to attitudinal measures, we collected and analyzed usage logs (anonymized by removing login and network address information) for actual behavior changes across the control and treatment groups. Table 2 shows the metrics derived from the usage log data.

To measure the treatment effect, we computed the ratio of the mean metric values (mean of observations for treatment vs. mean for control) within each of five demographic (sub)groups of the participants (Table 3). Significance (*p*-value) of each treatment/control ratio was tested using a resampling-based permutation test (e.g., see Hesterberg, Moore, Monaghan, Clipson, & Epstein, 2011, Sec. 16.5) with 10,000 resamples. (Such permutation tests are straightforwardly applied to statistics such as ratios and require fewer assumptions about the underlying distribution than do parametric tests.) Findings reported as significant remain significant with the false discovery rate (FDR) controlled at 0.05 ($m=63$, $\lambda=0.20$) (Storey & Tibshirani, 2003).

Caveats and Limitations

Two caveats relating to software bugs were identified and fixed in the trial's first two weeks. First, Google.com search results page costs were shown to both control and treatment groups during the first two weeks. These cost displays did positively influence the cost awareness of the control group (this effect is discussed in detail under "Findings"). Second, the billing system was set to price mobile data at a rate that was 1 Gp lower than originally intended (2 Gp/MB vs. 3 Gp/MB). While corrected, both bugs surely had *some* effect on user behavior. At the same time, they helped highlight aspects of Internet usage, such as how influential page costs were in driving cost awareness and how sensitive participants were (or became) to web pricing.

It is also important to note the limitations of our trial design. While we tried to include nonstudents, our trial largely focused on students because of the university environment. We placed limitations on the types of content accessible through the phones because of technical implementation and security constraints. Our participants could browse all web pages, but accessing apps, viewing videos, and downloading files were disabled. Apps and videos are not discrete resources and can continuously pull data; hence, size and length information cannot be known before streaming or downloading. Downloading files placed a security risk on the phones.

Findings

We now turn to our quantitative findings from logs analysis and qualitative findings from user interviews. The terms "SmartBrowse users" and "treatment users" are used interchangeably.

Simple satisfaction surveys (1 = extremely dissatisfied to 7 = extremely satisfied) showed that treatment

Table 2. Metrics Computed from Log Data.

Per-user metrics (control and treatment)	Fraction of days active	Frequency of user activity over the 10-week period
	Daily page views	Number of web pages viewed per day
	Daily sessions	Number of browsing sessions per day ("session" = continuous page views for 5-minutes)
	Session length	Number of pages per browsing session
	Page cost	Cost (GH₵) per page view
	Session cost	Cost (GH₵) per browsing session
	Total money spent	Total spend (GH₵) during the study
Per-user metrics (treatment only)	Total SmartAlerts seen	Total number of pages, SmartAlerts raised
	Total SmartAlert price	Total projected page costs, SmartAlerts seen
	Total SmartAlert savings	Total projected page costs, SmartAlerts declined
Per-website metrics	Number of visits	Used to compute popularity ranking during study

Table 3. (Sub)Groups Analyzed for Treatment Effect (see Table 1 for breakdowns).

(Sub)Group	Definition/Composition
All users	All study participants
Prior mobile data user	"prior" = had used mobile data prior to study "new" = had not used mobile data prior to study
Gender	female, male
Occupation	student, nonstudent
Income bracket	low, medium, high

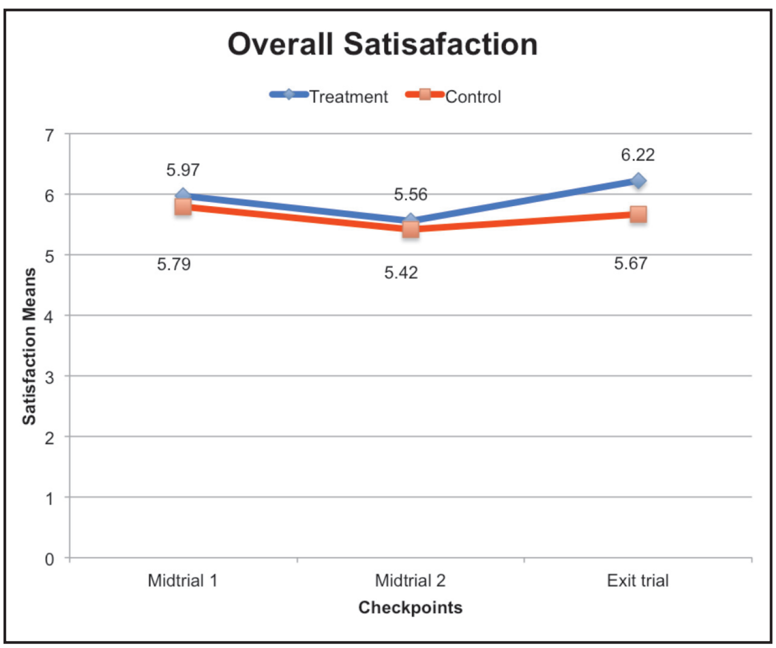

Figure 2. Overall satisfaction, measured across checkpoints.

users were marginally more satisfied than control users (Figure 2). We point this out to suggest that the experience that included the extra displays and interstitial SmartAlerts was not heavyweight or severe enough to cause undue dissatisfaction.

Mobile Data Usage Logs

SmartBrowse Users Went Online More Often Than Control Users

Over the 10-week period both control and treatment users, including new mobile data users, used their phones regularly without major dips in use (Figure 3).

We measured the number of users who effectively dropped out of the experiment by failing to use their phones regularly. To

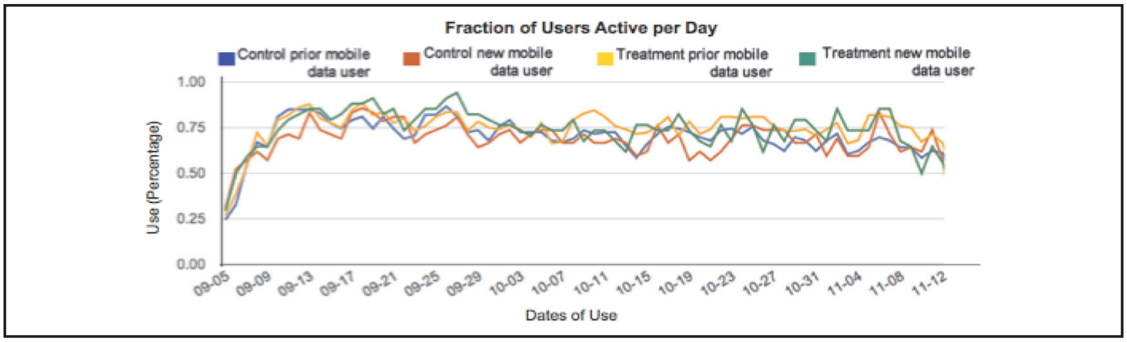

Figure 3. Fraction of users active per day over 10 weeks.

Table 4. Selected Per-User Metrics, Pooling Prior/New Mobile Data Users (metrics as defined in Table 2).

	Control (n = 148)	Treatment (n = 151)	Effect (%)	Significance (%)
Fraction of days active	0.704	0.753	7	99
Average daily page views	29.6	27.5	−7	not significant (69)
Average page cost (GH₵)	0.362	0.318	−12	95
Average session cost (GH₵)	0.0167	0.0132	−21	> 99
Money spent (GH₵)	5.07	4.11	−19	99

track drop-off rates, for each user we measured the fraction of days when they went online with their phone at least once (an "active day"). Over the 10-week experiment, on average, a user in the control group was active on 70% of the days, while a treatment user was active on 75% of the days. A treatment user was active an average of 7% more days than a control user (significant to the 99% level).

SmartBrowse Users Spent Less on Internet Credit

As seen in Table 4, users in the treatment group spent an average of 19% less than those in the control group during the trial (GH₵4.10 for treatment vs. GH₵5.06 for control; significant to the 99% level). While the magnitude of the treatment effect was direction₵lly consistent (i.e., −10–30%) across subgroups (as listed in Table 1), the effect in these smaller groups was not always significant.

Similar, significant reductions in average page cost and session cost can also be seen in Table 4. And as with total spend, these reductions were directionally consistent across subgroups but not at a significant level.

SmartBrowse Users Actively Responded to SmartAlert

About half of the money saved by treatment users is explained by users avoiding expensive pages via SmartAlert. (Recall that a SmartAlert is an interstitial warning page that asks a treatment user to confirm before loading an expensive page.) On average, the 148 control users spent GH₵450 during the 10-week trial, while the 151 treatment users spent GH₵360. As seen in Figure 4, the mode for all subgroups of treatment/control and prior/new mobile data users is in the GH₵3–4 range. When we total the costs of all pages that triggered SmartAlerts where the user chose Back (i.e., avoided loading the page and incurring the cost) instead of Continue (i.e., loaded the page and incurred the cost), this amounts to GH₵50, or about 55% of the difference between the two groups' total expenditure. In other words, treatment users potentially would have spent about GH₵50 more if we had simply taken away the SmartAlerts. While this is not a perfect conclusion (it is possible that users went back and loaded an alternative page instead, which we discuss in a later section), it is indicative of the potential benefit of SmartAlerts. Some treatment users noted that they had not realized the web pages they usually visited cost so much, motivating them to stop loading the web pages.

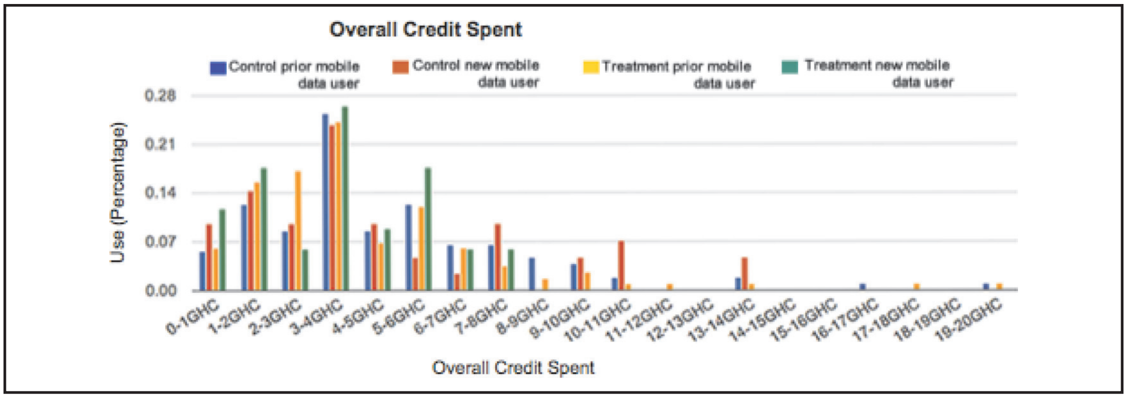

Figure 4. Distribution of total mobile data credit spend by individuals during the trial.

During the course of the experiment, treatment users saw about 2,000 SmartAlerts. On about 1,200 SmartAlerts, users chose Continue and paid to load the page. These accepted pages had an average cost of 4.5 Gp (maximum cost: 27 Gp). On 800 SmartAlerts, users chose Back or navigated elsewhere. These declined pages averaged 6 Gp (maximum cost: 42 Gp). On average, then, the declined pages were about 33% more expensive than accepted pages, showing that treatment users were indeed more likely to decline to pay for expensive pages.

SmartBrowse Users Went to Less Expensive Web Pages Overall

While Table 4 shows no significant difference in the average number of web pages visited per day, it also shows that treatment group users spent 12% less per page (significance 95%) and 21% less per session (significance > 99%) than control group users without compromising their overall information needs (more in the "Finding Less Expensive Alternatives" section). Recall from Table 2 that a session is defined as a series of web pages viewed without a > 5-minute break.

New Mobile Data SmartBrowse Users Experienced Greater Benefits

While we focused the trial on prior mobile data users, we were interested by the effect of SmartBrowse on new mobile data users (subjects who had not used the mobile Internet on their phones before the trial). As can be seen in Table 5, new users in the treatment group visited 53% more web pages per day than those in the control group (96% significance), with greater treatment reductions in spending at the page (−27%, 99% significance) and session (−34%, > 99% significance) levels than for the overall participant population.

From this, we *speculate* that cost-related information is more useful to new users, who have little or no internalized experience of the cost of browsing. This is intuitively supported by the much heavier tail in the spending distribution of new mobile data users in the control group vs. the new mobile data users in the treatment group (Figure 4). However, the reduction in total spend for new users (while greater than for the overall population, on average) was *not* significant, so this metric remains a topic for additional study.

Cost Awareness

SmartBrowse Users Had Greater Cost Awareness

When asked in the trial exit survey to assess whether their cost awareness had increased, just over half the control users answered affirmatively, whereas the vast majority of treatment users generally felt they had improved (57% of control participants vs. 92% of treatment participants). The self-assessment is supported by the cost awareness exercise results. As can be seen from the orange-highlighted mode values in Table 6, both control and treatment users came closer to the correct cost estimates as the trial progressed. Interestingly, most control participants seemed to have noticed the page cost labels on Google.com within the first two weeks—exposed to them by one of the previously noted bugs—and revised their estimates of how much web pages cost.

Table 5. Selected per-user metrics, for new users only (metrics as defined in Table 2)

	Control (n=42)	Treatment (n=34)	Effect (%)	Significance (%)
Fraction of days active	0.687	0.749	8	not significant (89)
Average daily page views	22.6	34.6	53	96
Average page cost (GH₵)	0.00364	0.00266	−27	99
Average session cost (GH₵)	0.0172	0.0114	−34	>99
Money spent (GH₵)	4.90	3.82	−22	n.s. (93)

Table 6. Cost Awareness Exercise Results. (Correct intervals are highlighted in green. The modes of each column are highlighted in orange.)
(a) "Text-lite" pages (on Google)

		Midtrial 1 (Week 3)		Midtrial 2 (Week 7)		Trial exit (Week 10)	
Est. cost (GH₵)	Baseline	Control	Treatment	Control	Treatment	Control	Treatment
< 0.01	3	18	6	9	10	13	10
0.01–0.05	36	61	66	80	115	94	133
0.06–0.15	60	22	27	25	13	16	8
0.16–0.50	111	25	20	25	4	16	6
0.51–1.00	25	10	4	2	0	0	0
1.01–5.00	28	3	6	0	0	0	0
> 5.00	4	2	0	0	0	0	1

(b) "Images + text" pages (on Facebook)

		Midtrial 1 (Week 3)		Midtrial 2 (Week 7)		Trial exit (Week 10)	
Est. cost (GH₵)	Baseline	Control	Treatment	Control	Treatment	Control	Treatment
< 0.01	7	12	9	4	13	13	12
0.01–0.05	16	41	59	66	106	84	111
0.06–0.15	45	23	19	24	24	23	18
0.16–0.50	141	44	40	21	9	14	12
0.51–1.00	32	15	8	0	1	0	0
1.01–5.00	28	4	0	0	0	0	0
> 5.00	2	0	3	1	0	0	0

(c) "Image heavy" pages (on Wikipedia)

		Midtrial 1 (Week 3)		Midtrial 2 (Week 7)		Trial exit (Week 10)	
Est. cost (GH₵)	Baseline	Control	Treatment	Control	Treatment	Control	Treatment
< 0.01	7	12	9	4	13	13	12
0.01–0.05	42	61	62	82	111	98	121
0.06–0.15	62	26	26	29	13	16	10
0.16–0.50	120	27	25	26	4	7	5
0.51–1.00	27	11	5	1	1	0	0
1.01–5.00	11	3	3	0	0	0	0
> 5.00	6	0	1	0	0	0	0

However, as can be seen from the green-highlighted (0.01–0.05 Gp) values in Table 6, more treatment users were in the correct cost estimate bracket than control users (both in proportional as well as absolute terms).

Treatment users' discussions of cost awareness often centered on the interstitial SmartAlerts. For example, several treatment users recalled events in which websites turned out to be more expensive than they originally assumed—one participant discovered that buybay.com cost him 46 Gp instead of 10 Gp as assumed. Conversely, some treatment users discovered that some favorite websites that they assumed were expensive did not cost so much. For example, one user said,

> I used to think Facebook was very costly, especially for photos. After I started SmartBrowse, I realize it is not costing so much. I am on Facebook always and it costs only 1–2 Gp.

SmartBrowse Users Reported Being Better at Data Credit Management

In the trial exit survey, 72% of treatment users claimed that they were making better decisions at managing their Internet credit (the remainder reported no improvement). One of those users said,

> We don't have much pocket money. I buy GH₵1–2 credit from our Hall shop every 2–3 days. Earlier I did not notice, but now I know which websites are cheap and my balance is OK. It lasts 4–5 days now.

By contrast, only 43% of control users felt they were making better decisions.

Feeling in Control with Mobile Operators

Many treatment users mentioned they felt they were in better control of mobile operator pricing after using SmartBrowse. By learning the cost of web pages and setting their expectations around what various types of content should cost, users mentioned they had a better grasp of prices and that experience could be applied when purchasing from mobile operators. One user said,

> Now I know how much to pay for which website. I will not get cheated by any phone company in the future.

User Strategies for Saving Credit by Becoming Cost Conscious

Several participants in both the treatment and control groups stated that they spent more time online since starting the SmartBrowse trial. Many participants reported going online more during the first two weeks of the trial to explore the phone, visit various websites, and get a sense of the pricing of SmartBrowse. One participant said,

> Wednesday to Monday, I spent my first GH₵2. I finished the second GH₵2, which I bought, in a week. I used my own money in the second week, so I was careful with browsing. First GH₵2 was an incentive to use, so I visited pages I was not supposed to visit.

Several treatment users reported becoming more aware of their mobile data spending since starting the SmartBrowse trial. The visual foregrounding of usage information by SmartBrowse was cited as a huge influence in becoming more spending-conscious. For example, Ghanaleaks.com amount of credit is so high. Within 2–3 minutes, about 80 Gp-GH₵1 will be gone. I did not know it costs so much. Now I know. I don't open that site anymore.

Finding Less Expensive Website Alternatives

With SmartBrowse, many prior mobile data users in the treatment group found less expensive alternatives to the websites they wanted to visit by finding substitutes in the same content category. For example, one participant noted,

> I used to go to soccernet.com every day to check football scores. But then I realized it costs 18Gp from the pop-up [SmartAlert]. So I stopped going there and now I go to goal.com instead, which costs only 5–6 Gp.

Many treatment users mentioned they would visit a website regardless of cost if it provided the information they wanted, such as Facebook and MyJoyOnline, without compromising their information needs. The most-visited websites (e.g., Facebook, Google, Ghanaweb, MyJoyOnline, Twitter) redirected to their mobile versions automatically. When users were less particular about visiting a specific website (e.g., dictionary sites or research sites), a less expensive option was chosen. For content downloading and video consumption needs, participants reported using other devices, and this behavior may have implications for cost transparency tools in not restricting users' browsing freedom or agency.

New mobile data users did not have a strong conception of go-to websites on the phone and were generally more cost conscious in their web browsing decisions.

Using SmartBrowse Features

Balance Bar

The balance bar was generally perceived to be valuable. For many treatment participants, the balance bar served as a way to monitor their spending, both in aggregate terms and for individual pages. One said,

> Before I start browsing, I have to know amount of credit on the phone. So I will check my balance on the bar. Within five minutes, I will check the credit deducted. Once in a while to see if credit is OK, to make sure it is not jumping from GH₵2 to GH₵1.50 suddenly.

Top-Up Page

Control group participants found the visual display of balance on the top-up page to be convenient and simple. The top-up page was seen as more reliable than USSD balance checking, which frequently resulted in "network error" and "timeout" messages. Many users liked the ease of inputting the short scratch code. Several participants mentioned that they checked the top-up page frequently initially but stopped checking as much once they started trusting the system, feeling "safe" with SmartBrowse charges and controls.

SmartAlert

SmartAlert seemed less useful when a decision to visit a website was already made. A few treatment users mentioned they grew annoyed when SmartAlert popped up over their favorite websites. One user said,

> Cosmo is my favorite site! I have to check it every day. But the alert pops up every time I visit it. It's quite annoying.

However, users were reluctant to change the SmartAlert settings to a lower threshold (among those who had discovered the link), fearing that there would not be an alert for expensive but "fun" web pages that they would not necessarily visit if they knew the cost. A system to adaptively adjust the threshold could mitigate this issue.

Web Page Costs

The web page cost displays on the Google.com search results pages were perceived as "good to know." Decisions on which Google.com search results to visit were highly contextual, i.e., content influenced most visit decisions, and cost influenced some decisions. As such, web page cost information was largely applied in deciding to visit educational websites (less so for entertainment websites).

New Behaviors and Changes

Shortcuts Were Key to Defining Browsing Behavior

The shortcuts we placed on the Android home screen introduced new browsing patterns to many participants. Several noted that they had visited unfamiliar websites simply because there was a shortcut to the site, such as MyJoyOnline.com and Goal.com. In fact, most websites visited by participants seem to be within the scope of the shortcuts (not surprising, given that the shortcuts pointed to Ghana's most-visited websites).

Going Online on a Phone for the First Time

In the case of new mobile data users who had some prior Internet experience with cybercafés and PCs, many participants reported that they had enjoyed the smartphone Internet experience and that it had changed their perception of the difficulty of browsing on phones. Some participants mentioned they had gained an awareness of web page costs in general, which they had not consciously considered before:

> Before SmartBrowse, I did not think much about how much websites cost. Now I know. I will avoid some sites, even in café.

Of course, page sizes and data costs vary across different forms of access—mobile vs. wired, PC vs. phone, etc.

New mobile data users with no prior Internet experience had some difficulty using the Internet because of literacy constraints. Some participants reported they sought the help of friends and family members to browse.

Getting Used to a Smartphone

All interviewed participants responded positively to using the phone, with varying levels of enthusiasm depending on whether they were prior smartphone users. New smartphone users identified the phone qualities—touch screen, large screen size, ability to zoom ("It has life!" in the words of one participant), and portability—as highly enjoyable and convenient. Prior smartphone users were less enthusiastic about the Motorola dual-SIM phone. New mobile data users reported enjoying not having to visit a cybercafé and being able to instantly retrieve information on the Internet anytime, anyplace.

Participants paid a lot of attention to network speeds, looking for network strength bars and page loading speed, which appeared to be a major determinant of satisfaction.

Experiences After the SmartBrowse Trial

As mentioned earlier, 42% of participants (n = 126) responded to the post-trial follow-up survey. Roughly half the respondents stated they had purchased a new phone since the end of the trial. Of these, 86% from the control group and 78% from the treatment group stated they had purchased a smartphone. Most upgrades were from midrange feature phones to smartphones. Half the new smartphone purchasers attributed their decision to purchase a smartphone to the SmartBrowse experience. Of the respondents, 36% had converted to regular mobile data users since the trial (54% had been such prior to the trial; only 10% did not convert). Roughly half had moved to the mobile phone as their primary Internet browsing device.

We conducted a cost estimation exercise in the follow-up survey. Both control and treatment participants were mostly accurate in their estimates, possibly suggesting building a *knowledge* of web page costs that persisted beyond the trial. Participants noted that SmartBrowse changed their Internet use by making them cost conscious when using the Internet and made them appreciate browsing the Internet on the phone. One said,

> I now know which site to visit and not to visit. I now spend little time to browse cos (because) of cost involved. That is, I try to save credit.

Another said,

> [I changed my Internet use] because am very much conscious with page charge when using a friend's phone to browse.

Discussion

What broader learnings from our price transparency trial can be applied to ICTD?

Empowerment Through Usage Transparency

While not definitive, our study is a positive indication that price and usage transparency, when applied to an opaque and expensive resource, can bring about economic efficiency. By providing a balance of usage information with actionable prompts before appropriate spending thresholds, SmartBrowse created a sense of empowerment among users by helping them control their spending and, perhaps, make better-informed choices about mobile operators.

On the other hand, transparency of price and use leads users to become mindful of web page costs, with the potential for negative influence on their decisions to visit certain websites. While our trial showed no specific evidence of lessening user satisfaction, the trial was limited to a student population and a relatively short period of three months. Longer-term and lab studies with other user groups could lead to a deeper understanding of the implications of behavior change through price transparency and assess the degree to which cost considerations lead to a suboptimal level of satisfice.

Easing in the Novice User's Experience of Mobile Data

With new mobile data users, we found that price transparency led to a better understanding of the Internet, which had been perceived as expensive and beyond the reach (among other perceptions) of many low-income mobile users. Price transparency allowed new mobile data users to get what they wanted from the Internet, both by scaffolding their cost understanding of the Internet and by creating a feeling of control. With friends and family around them using mobile data price transparency tools, new mobile data users may be introduced to an Internet with the expectation of being able to manage their experience and expenditures.

The Emerging Smartphone Users

Our study points to high-tech interventions with urban emerging regions' users as a promising area of ICTD research. With rising rural-to-urban migration rates in many emerging regions, the urban low-income groups are increasingly interfacing with similar technologies as other income brackets. In particular, dropping prices, extensive distribution channels, and increased familiarity with technology have led growing numbers of urban consumers to adopt high-tech ICTs such as smartphones. At the time of this trial, Informa estimated smartphone penetration in Africa at 11% (79M (million) of 742M connections worldwide); corresponding estimates for 2015 and 2018 are 20% (204M of 1,012M) and 34% (412M of 1,213M), respectively (Informa, 2014). As smartphone prices drop to the US$25 point, experiences of the kind we present here will quickly become more representative of what middle- and lower-income groups experience.

ICTD has predominantly focused on rural socioeconomic groups, but these emerging urban segments are not only relevant to cities, but are also linked in complex ways to rural settlements. Families in sub-Saharan Africa are often multispatial, with household members who diversify income by moving and separating as opportunities arise and conditions change (Ackah & Medvedev, 2012; van Dijk, Foeken, & van Til, 2001); internal migration has led to increased economic and social ties back to villages in the form of remittances and of circulation of mobile phones, often as gifts (Sey, 2011). Even rural populations in emerging regions will have increasing exposure to mobile data and smartphones due to circulation.

The Promise of Phone Trials

As evidenced in our follow-up survey, a forward-looking technology trial can heavily influence future purchase and usage decisions, provided the devices are within economic and physical reach of the user. We note that deploying with students could have positively influenced smartphone uptake. More research on smartphone trials can help us understand what skills and capabilities smartphones engender and what drawbacks they present. Smartphone trials may provide a great mechanism for introducing mobile data in new communities. Well-thought-out shortcuts may steer traffic toward development content, such as Khan Academy or Wikipedia.

Forward-Looking Research in ICTD

Our trial presents forward-looking research wherein users are exposed to the next generation of technologies. Traditionally, ICTD research has focused on existing methods such as voice calls and SMS for cost, scalability, and implementation reasons. As we have shown, it is informative to conduct forward-looking trials to provide a glimpse into the problems that low- and middle-income users may encounter in the near future.

On the other hand, research with new technology involves practical challenges that must be planned for. Distributing expensive devices (the phones in the trial cost US$150 each, more than the average monthly pocket money of student participants) in longitudinally for 10 weeks involves its own set of challenges: Users may need training, they may need help solving technical issues, and phones may be lost.

In the next section, we describe how we set up the trial, trained participants, and designed appropriate incentives.

Running the SmartBrowse Trial

Trial Preparation

The trial was conducted at the University of Ghana, Legon, a few kilometers outside central Accra. The university environment enabled us to connect participants' ID cards to their Android phones for the 10-week trial period, although no action was taken when phones were actually reported lost during the trial.

Our Ghana ground team comprised a research coordinator, a research assistant, and six interns who helped run the study. A dedicated room, called the SmartBrowse Hub, was set up in a central location. Interns were constantly available to troubleshoot phone issues, sell scratch cards, and answer questions. The Hub was instrumental in making the trial work, allowing users to walk in with issues on their way to classes.

Motorola dual-SIM XT685 Android smartphones were provided to all participants, preloaded with an MTN SIM card exclusively provisioned for mobile data. Stickers were placed on each phone's back plate with the ground team's phone number for assistance. Custom scratch cards were designed for use with the SmartBrowse top-up page in denominations of GH₵1, GH₵2, and GH₵5

Each participant received a bag with the phone, a SIM, a charger, an instruction booklet for the phone, a campus map with directions to the Hub, and a free starter GH₵2 scratch card.

Trial Recruitment

The trial recruitment was broken into a series of steps: *invitation,* where interns approached students in high-traffic areas of the university; *ID check,* where the candidate ID was checked to verify university affiliation; *screening,* where the recruited candidates were screened based on criteria listed in the "Research Design" section and assigned to control or treatment groups; *paperwork* and *orientation,* where participants were required to complete a baseline survey and shown how to use the phone including setup, credit top-up, making phone calls, and loading web pages.

Incentives

We provided a monetary incentive of GH₵200 (roughly US$100) to each participant, distributed in parts at the three trial checkpoints. We set the two midtrial incentives (GH₵30), high enough to motivate the participants to travel to the Hub (both to ensure we were collecting survey data from participants each month and to ensure the phones were working correctly). We made the final incentive payment (GH₵140), high enough to motivate our participants to return the phone at the end of the trial.

Conclusion

In this article, we presented our findings from a 10-week study of a mobile data price transparency tool with 299 participants. Our findings pointed to a net increase in Internet use with a net decrease in credit consumption as well as greater cost awareness among the treatment users. We discussed the logistics of running such a trial longitudinally in an emerging region context.

Price transparency is important to informed access and Internet use. As lower-income user groups typically encounter the Internet for the first time on mobile phones, providing ways to be aware of and control their data expenditures is important for economic efficiency. While our trial focused on mobile data spending on smartphones, the techniques can be easily applied to bandwidth management in general. Conducting forward-looking research into technologies that are slowly but steadily flowing into emerging regions can help preemptively reduce barriers to technology use and help millions of new and existing users experience a safer, more manageable, less erratic Internet. ■

Acknowledgments

Our sincere gratitude to Pedel Oppong, Diana Akrong, Godfred Antwi, Louis-Mark Lartey, Benjamin Odame, Daisy Prempeh, and Janet Teng for their dedication and hard work in making the trial successful. We also thank Ama Dadson, Dr. Robert Sowah, Angela Crandall, and Jessica Colaço for their help with the research.

Nithya Sambasivan, User Experience Researcher, Google Inc., USA. nithyasamba@google.com

Paul Lee, Senior Product Manager, Google Inc., USA. paullee@google.com

Greg Hecht, Staff Software Engineer, Google Inc., USA. ghecht@google.com

Paul M. Aoki, Software Engineer, Google Inc., USA. aoki@acm.org

Maria-Ines Carrera, Software Engineer, Google Inc., USA. mariaines@google.com

Jenny Chen, Software Engineer, Google Inc., USA. jennyc@google.com

David Cohn, Director, Natembea Foundation, USA. david.cohn@gmail.com

Pete Kruskall, CTO & Cofounder, Kensho, USA. pete.kruskall@gmail.com

Everett Wetchler, Data Scientist, Bayes Impact, USA. everett.wetchler@gmail.com

Michael Youssefmir, Access Technology Development Principal, Google Inc., USA. myoussefmir@google.com

Astrid Twenebowa Larssen, Designer, Halogen AS, Norway. twenebowa@gmail.com

References

Ackah, C., & Medvedev, D. (2012). Internal migration in Ghana: Determinants and welfare impacts. *International Journal of Social Economics, 39*(9–10), 764–784. doi:10.1108/03068291211253386

Anokwa, Y., Ribeka, N., Parikh, T., Borriello, G., & Were, M. C. (2012). Design of a phone-based clinical decision support system for resource-limited settings. In *Proceedings of ICTD 2012* (pp. 13–24). New York, NY: ACM. doi:10.1145/2160673.2160676

Butkiewicz, M., Madhyastha, H., & Sekar, V. (n.d.). Understanding website complexity: Measurements, metrics, and implications. In *Proceedings of the 2011 ACM SIGCOMM Conference on Internet Measurement* (pp. 313–328). New York, NY: ACM. Retrieved from http://goo.gl/3GFwP2

Chabossou, A., Stork, C., Stork, M., & Zahonogo, Z. (2009). *Mobile telephony access and usage in Africa* (pp. 392–405). New York, NY: IEEE. doi:10.1109/ICTD.2009.5426688

Chava, S., Ennaji, R., Chen, J., & Subramanian, L. (2012). Cost-aware mobile web browsing. *IEEE Pervasive Computing, 11*(3), 34–42. doi:10.1109/MPRV.2012.19

Chen, J., & Subramanian, L. (2013). *Interactive web caching for slow or intermittent networks* (pp. 1–10). New York, NY: ACM Press. doi:10.1145/2537052.2537057

Chetty, M., Banks, R., Brush, A. J., Donner, J., & Grinter, R. (2012). "You're capped!" Understanding the effects of bandwidth caps on broadband use in the home. In *Proceedings of CHI 2012* (pp. 3021–3030). New York, NY: ACM. doi:10.1145/2207676.2208714

Chetty, M., Banks, R., Harper, R., Regan, T., Sellen, A., Gkantsidis, C., . . . Key, P. (2010). Who's hogging the bandwidth?: The consequences of revealing the invisible in the home. In *Proceedings of CHI 2010* (pp. 659–668). New York, NY: ACM. doi:10.1145/1753326.1753423

Communications Chambers. (2012). *Broadband pricing database—Explanatory notes.* Retrieved from http://goo.gl/5mxCkU

Crandall, A., Otieno, A., Mutuku, L., Colaço, J., Grosskurth, J., & Otieno, P. (2012, November). *Mobile phone usage at the Kenyan base of the pyramid.* Washington, DC: *info*Dev. Retrieved from http://goo.gl/sOOhi

Donovan, K., & Donner, J. (2010). A note on the availability (and importance) of pre-paid mobile data in Africa. In *Proceedings of M4D 2010* (pp. 263–267). Retrieved from http://goo.gl/w4GqJs

Froehlich, J., Findlater, L., & Landay, J. (2010). The design of eco-feedback technology. In *Proceedings of CHI 2010* (pp. 1999–2008). New York, NY: ACM. doi:10.1145/1753326.1753629

Gillet, J., Hatt, T., & Lucini, B. A. (2013). *Tailoring mobile Internet tariffs for prepaid users—A balancing act.* GSMA. Retrieved from http://goo.gl/S8ZqTA

Haughton, J. H., & Khandker, S. R. (2009). Measuring poverty. In *Handbook on poverty and inequality* (pp. 9–37). Washington, DC: World Bank. Retrieved from http://goo.gl/SMu8MR

Hesterberg, T., Moore, D. S., Monaghan, S., Clipson, A., & Epstein, R. (2011). Bootstrap methods and permutation tests. In D. S. Moore & G. P. McCabe, *Introduction to the practice of statistics* (7th ed.). New York, NY: W. H. Freeman. Retrieved from http://goo.gl/BGOW0f

iHub Research. (2012a). *Mobile Internet usage field survey: Ghana.* Nairobi, Kenya: iHub Research. Retrieved from http://goo.gl/xJlAkj

iHub Research. (2012b). *Mobile Internet usage field survey: Kenya, Nairobi.* Nairobi, Kenya: iHub Research. Retrieved from http://goo.gl/xIo7oD

Informa Telecoms & Media. (2014). *Africa telecoms outlook 2014: Maximizing digital service opportunities.* Retrieved from http://goo.gl/nt4CNM

International Telecommunication Union (ITU). (2012). *Measuring the information society.* Geneva, Switzerland: ITU. Retrieved from http://goo.gl/agHQWi

International Telecommunication Union (ITU). (2013). *Measuring the information society.* Geneva, Switzerland: ITU. Retrieved from http://goo.gl/vrnaId

Isaacman, S., & Martonosi, M. (2011). Low-infrastructure methods to improve Internet access for mobile users in emerging regions. In *Proceedings of WWW 2011* (pp. 473–482). New York, NY: ACM. doi:10.1145/1963192.1963361

Minges, M. (1999). Mobile cellular communications in the Southern African region. *Telecommunications Policy, 23*(7–8), 585–593. doi:10.1016/S0308-5961(99)00039-7

Moore, J. C., Stinson, L. L., & Welniak, E. J. (2000). Income measurement error in surveys: A review. *Journal of Official Statistics, 16*(4), 331–361. Retrieved from http://goo.gl/YOxHMj

Murphy, L. L., & Priebe, A. E. (2011). "My co-wife can borrow my mobile phone!": Gendered geographies of cell phone usage and significance for rural Kenyans. *Gender, Technology and Development, 15*(1), 1–23. doi:10.1177/097185241101500101

Narayanan, S., Chintagunta, P. K., & Miravete, E. J. (2007). The role of self selection, usage uncertainty and learning in the demand for local telephone service. *Quantitative Marketing and Economics, 5*(1), 1–34. doi:10.1007/s11129-006-9015-z

Otsuka, T. (2009). *Mobile broadband: Pricing and services.* Paris, France: Organisation for Economic Co-operation and Development. Retrieved from http://goo.gl/GDdbXs

Pentland, A., Fletcher, R., & Hasson, A. (2004). DakNet: Rethinking connectivity in developing nations. *Computer, 37*(1), 78–83. doi:10.1109/MC.2004.1260729

Prahalad, C. K. (2005). *The fortune at the bottom of the pyramid: Eradicating poverty through profit.* Upper Saddle River, NJ: Wharton School Publishing. doi:10.1111/j.1468-0270.2008.864_4.x

Ramachandran, D., Canny, J., Das, P. D., & Cutrell, E. (2010). Mobile-izing health workers in rural India. In *Proceedings of CHI 2010* (pp. 1889–1898). New York, NY: ACM. doi:10.1145/1753326.1753610

Sambasivan, N., & Cutrell, E. (2012). Understanding negotiation in airtime sharing in low-income microenterprises. In *Proceedings of CHI 2012* (pp. 791–800). New York, NY: ACM. doi:10.1145/2207676.2207791

Sambasivan, N., & Smyth, T. (2010). The human infrastructure of ICTD. In *Proceedings of ICTD 2010* (pp. 1–9). New York, NY: ACM. doi:10.1145/2369220.2369258

Sey, A. (2011). "We use it different, different": Making sense of trends in mobile phone use in Ghana. *New Media & Society, 13*(3), 375–390. doi:10.1177/1461444810393907

Storey, J. D., & Tibshirani, R. (2003). Statistical significance for genomewide studies. In *Proceedings of the National Academy of Sciences, 100*(16), 9440–9445. doi:10.1073/pnas.1530509100

Stork, C., Calandro, E., & Gillwald, A. (2013). Internet going mobile: Internet access and use in 11 African countries. *Info, 15*(5), 34–51. doi:10.1108/info-05-2013-0026

Sun Basorun, A., Moodley, L., Moraje, S., & Nielsen, M. (2012). *Understanding the African digital consumer.* New York, NY: McKinsey & Co. Retrieved from http://goo.gl/q6smk

Van Dijk, H., Foeken, D., & van Til, K. (2001). Population mobility in Africa: An overview. In M. de Bruijn, R. van Dijk, & D. Foeken (Eds.), *Mobile Africa: Changing patterns of movement in Africa and beyond.* Leiden, The Netherlands & Boston, MA: Brill. Retrieved from http://goo.gl/E3jl9R

Vodacom. (2012). *Vodacom Group Limited integrated report.* Retrieved from http://goo.gl/ei1WF

von Tetzchner, J. S. (2010, March 29). *State of the mobile web, February 2010.* Oslo, Norway: Opera Software ASA. Retrieved from http://goo.gl/RLuYcJ

Walton, M., & Donner, J. (2012). *Public access, private mobile: The interplay of shared access and the mobile Internet for teenagers in Cape Town.* Cape Town, South Africa: University of Cape Town. Retrieved from http://goo.gl/L9xct

Wang, X. S., Balasubramanian, A., Krishnamurthy, A., & Wetherall, D. (2013). Demystifying page load performance with WProf. In *Proceedings of the 10th USENIX Conference on Networked Systems Design and Implementation* (pp. 473–485). Berkeley, CA: USENIX Association. Retrieved from http://goo.gl/a7IgYt

Wang, X. S., Shen, H., & Wetherall, D. (2013). Accelerating the mobile web with selective offloading. In *Proceedings of the 2nd ACM SIGCOMM Workshop on Mobile Cloud Computing* (pp. 45–50). New York, NY: ACM. Retrieved from http://goo.gl/Beq0ig

Williams, M. D. J., Mayer, R., & Minges, M. (2011). *Africa's ICT infrastructure: Building on the mobile revolution.* Washington, DC: World Bank. Retrieved from http://goo.gl/6uHMOl

Wyche, S. P., Smyth, T. N., Chetty, M., Aoki, P. M., & Grinter, R. E. (2010). Deliberate interactions: Characterizing technology use in Nairobi, Kenya. In *Proceedings of CHI 2010* (pp. 2593–2602). New York, NY: ACM. doi:10.1145/1753326.1753719