Journal of Development Economics 154 (2022) 102762

Contents lists available at ScienceDirect

Journal of Development Economics

journal homepage: www.elsevier.com/locate/devec

Regular article

Network adoption subsidies: A digital evaluation of a rural mobile phone program in Rwanda☆

Daniel Björkegren[a,*], Burak Ceyhun Karaca[b]

[a] Brown University, United States of America
[b] Keystone Strategy, United States of America

ARTICLE INFO

JEL classification:
O33
O14
L96

Keywords:
Targeting
Subsidies
Policy audits
Mobile phones
Networks
Technology adoption

ABSTRACT

Governments commonly target subsidies for communication networks in remote areas. We analyze a program in Rwanda that subsidized the equivalent of 8% of the stock of mobile phones, using 5.3 billion transaction records. Handsets mostly stayed in rural areas: 85% of accounts receiving subsidized handsets mostly use a rural tower. Subsidized handsets were used as much as purchased handsets. Recipients are highly connected to each other. We simulate welfare effects using a network demand system. Up to 69% of the impact on operator revenue comes from spillovers on nonrecipients. We also assess counterfactual targeting based on network properties and vouchers.

1. Introduction

Governments and NGOs have spent billions of dollars to connect poor and rural consumers to information and communication networks (Garbacz and Thompson Jr., 2005; GSMA, 2013). One common intervention is to subsidize access points, which can generate spillover benefits across the social network.[1] How should adoption subsidies be targeted?

This paper studies the targeting of an mobile phone handset adoption program implemented by the Rwandan government, using data from 5.3 billion transaction records from Rwanda's dominant mobile phone operator. The program allocated 53,352 handsets to rural areas, 8% of the stock of handsets at the time in 2008. We describe the implemented program using household surveys and digital data, use the estimated model of Björkegren (2019) to evaluate its welfare impacts, and then use that model to evaluate alternate network targeting rules.

The paper proceeds in four parts:

First, we analyze the implemented program using household surveys. Handsets were allocated to prioritized rural districts. Within these districts, local governments handled their own distribution, generally by allowing residents to self register for a subsidized phone ('decentralized distribution with self-targeting'). We assess the relationship between subsidy allocations and overall adoption, using household surveys aggregated to the district level. Being allocated subsidized handsets for an additional percentage point of households is associated with an increase of adoption between 1.88 and 3.31 percentage points between 2005 and 2010. This association exceeds one, which is consistent with adoption spillovers, but it is likely confounded by other factors.

Second, we provide descriptive evidence on the implemented program using digital data. Digital transactions reveal that most accounts were used in the rural areas that were prioritized, though there was some leakage (81% of subsidy recipients activated their accounts at a rural cell tower; after activation, 85% of recipients had a rural modal cell tower). The ultimate recipients of the handsets used them in a similar manner to individuals who paid market prices for their phones, suggesting the handsets were not wasted. Like other network targeting problems, the effects of the program depend on how recipients' adoption affected the entire network of potential adopters. Although

☆ Björkegren thanks Michael Kremer, Greg Lewis, and Ariel Pakes for guidance and encouragement, and Sendhil Mullainathan for helpful conversations. Thanks to Nathan Eagle for initial access to data, and the staff of the telecom partner, RURA, RDB, CGIS-NUR, NISR, Minicom, PSF, EWSA, and BRD for helpful conversations and access to data. This work was supported by the Stanford Institute for Economic Policy Research through the Shultz Fellowship in Economic Policy, and the JP Morgan Chase Research Assistant Program at Brown University.

* Corresponding author.
E-mail addresses: dan@bjorkegren.com (D. Björkegren), karacaburakceyhun@gmail.com (B.C. Karaca).

[1] For example, the Indian state of Chhattisgarh is implementing a program to provide smartphones to 5 million people (Ghose, 2018).

https://doi.org/10.1016/j.jdeveco.2021.102762
Received 7 October 2020; Received in revised form 6 October 2021; Accepted 15 October 2021
Available online 29 October 2021

we observe only people who eventually adopted, our data cover a period of exponential organic growth, so this includes many people who could have been induced to adopt earlier. Social networks are similar between those who ultimately used subsidized handsets and those who purchased phones at retail, suggesting that the handsets were not targeted to optimize network spillovers. Additionally, the dispersal program was relatively clustered: subsidy recipients are nearly 5 times more connected to each other than to the rest of the network. This could arise from natural clusters in rural areas, or from the program being shared by word of mouth.

In the third part of the paper, we evaluate the welfare effects of the implemented adoption subsidies across the network using the model and estimates of Björkegren (2019). That paper overcomes simultaneity in consumer adoption decisions by inferring the value generated by each connection from subsequent interaction across that connection. Calls are billed by the second, so a subscriber must value a connection at least as much as the cost of calls placed across it. Variation in prices and coverage identifies the underlying utility of communication across each link. Consumers choose when to adopt, by weighing the increasing stream of utility from communicating with the network against the declining cost of handsets. This allows us to back out bounds on each consumer's idiosyncratic value of having a phone. In equilibrium, each recipient reacts directly to the subsidy, and then each individual in the network reacts to each other's responses, capturing effects that ripple through the network. There tend to be multiple equilibria because individuals may coordinate on being optimistic or pessimistic about others' adoption; we identify the highest and lowest equilibria using the lattice structure of equilibria.

The decision to adopt a subsidized good only loosely identifies how much it is valued; as a result, bounds on impacts are wide and admit the possibility of zero impact in the upper bound of the highest equilibrium. However, in the lowest equilibrium and a focal equilibrium in which subsidy recipients value phones no more than retail adopters, welfare impacts are positive, with a social rate of return of at least 44%. A substantial portion of impact arises from spillovers. Nonrecipients account for up to 65% of the effect on revenue, and up to 24% of the effect on revenue arises from network ripple effects as nonrecipients change their adoption decisions.

In the final part of the paper, we use the same model to compare the performance of the implemented program (decentralized distribution with self-targeting) against alternatives that use full network information, and heuristics that can be implemented without network information.

We find there is a tradeoff to targeting based on centrality (Banerjee et al., 2013; Beaman et al., 2021): when subsidies are targeted to central nodes, 73% of consumer surplus accrues to urban consumers, rather than the remote, rural people this program aimed to benefit. While this could partially be adjusted by stratifying by location, there appears to be an efficiency/equity tradeoff: the impact on a simple sum of aggregate welfare is 32% lower when allocating to random rural nodes than to random urban nodes. This suggests that governments and nonprofits should consider whether their aims are simply to maximize diffusion, even if that means benefiting the best connected people, or whether they place higher welfare weights on individuals who are remote, either geographically or socially.

We also assess heuristics that can be implemented by policymakers who lack network information. We consider a common strategy used by growing tech companies, providing existing subscribers with a voucher that can be passed along to their strongest unsubscribed contact (like nomination or 'one hop' targeting strategies: Kim et al., 2015; Chin et al., 2021). This strategy exploits private information that individuals have about the people around them. In a focal equilibrium, vouchers would have increased aggregate welfare, by identifying individuals who will heavily use phones.

One caveat is that we only observe calls among people who eventually adopted, which trace out a partial social network. Because we study a period with substantial adoption, we can retrospectively analyze what would have happened to those adopters had they received a subsidy, but this approach could not be used to study the adoption of such a good prospectively in a setting with low adoption. However, we expect that social networks as revealed by digital interactions may be helpful for understanding spillovers in other goods or forms of information.

Related work

This paper demonstrates a new approach of digital policy evaluation, which uses streams of digital data to evaluate policies and programs. An increasing array of goods can be monitored digitally. This paper considers a good that automatically recorded a rich stream of data about how it was ultimately used, by whom, in what location of the country and in what location in the social network. These measures are taken continuously over the lifetime of the good at close to zero cost. This paper is a proof of concept that demonstrates how these measures could be used for a static policy evaluation; the same approach could be used for real time monitoring, allowing a way to adjust policies based on high frequency feedback.

We turn this approach towards an important policy question for developing countries. Digital connectivity can have substantial effects on developing countries (Jensen, 2007; Jack and Suri, 2014; Hjort and Poulsen, 2019); however, there is little empirical work on the economics of providing access to these networks.[2] Based on aggregate price elasticities, Garbacz and Thompson Jr. (2005) suggest handset subsidies could be beneficial. But operator groups suggest that universal access funds are not well spent and that goals would be better served by eliminating them and lowering taxes (GSMA, 2013). The last two sections of this paper use the structural model of mobile phone adoption in Rwanda from Björkegren (2019), which in that paper is used to assess the impact of nationwide taxes, and rural tower construction. We use this model to evaluate the impact of subsidy programs that are targeted to particular individuals in the network, and to compare the simulated outcomes of alternative subsidies.

Governments and NGOs distribute substantial resources through subsidies and in kind transfers. We build on a large literature that considers the structure of distributions: whether they reach intended recipients (Alatas et al., 2012, 2016), whether they are used (Cohen and Dupas, 2010), and how they compare to cash transfers (Cunha et al., 2019). But much of what we know about these programs comes from traditional administrative records, which may be coarse or not reflect on-the-ground realities, or surveys, which are costly to run and thus typically reach only small samples at discrete points in time.

In many contexts, the economic justification for subsidies is based on economic spillovers. Like many real world settings, our network is a single, interconnected network, and spillovers are diffuse. Diffuse spillovers are difficult to identify with traditional methods: a person may adopt after a contact adopts because the contact provides network benefits, or because they have similar traits or are exposed to similar environments. Randomized controlled trials commonly restrict attention to networks that have relatively disconnected subgraphs that can be independently randomized (Banerjee et al., 2013), or require careful corrections for interference (Athey et al., 2018). Alternately, the literature has restricted focus to sharply implemented quasiexperiments (for example, Higgins (2019) identifies spillover effects on retailers and other consumers from a large debit card rollout for government beneficiaries). The structural model of Björkegren (2019) allows us to trace the diffuse spillovers generated by the program as they ripple through the network of potential adopters.

[2] There is a literature on universal access in the U.S.: for example, the U.S. subsidized telephone access to the poor through the LifeLine program, with limited takeup and effects (Burton et al., 2007; Garbacz and Thompson, 1997).

Table 1
Allocation of subsidized handsets by district.

Household properties		District mean		Difference
		Participating	Nonparticipating	p-value
Rural		0.94	0.73	0.04
Consumption per capita		$204	$334	0.03
Handsets allocated	Total	3556.8	0	0.00
	Per household	0.05	0.00	0.00
Own mobile phone	2005	0.04	0.12	0.07
	2010	0.40	0.47	0.17
	Difference	0.36	0.36	0.76
N		15	15	

Source: Handset allocations: Banque Rwandaise de Developpement; other columns: EICV 2 and 3 surveys, National Institute of Statistics, 2005–2006, 2010–2011.

The last part of the paper connects to a literature that considers how information or influence diffuses through a network (Domingos and Richardson, 2001; Kempe et al., 2003). Several empirical studies have found that using social network information to target interventions can improve overall diffusion (Banerjee et al., 2013; Kim et al., 2015; Beaman et al., 2021). In developing country contexts, researchers have obtained social network information via survey measures, but this is costly (Perkins et al., 2015), so researchers have sought alternate methods such as using heuristics (Kim et al., 2015; Banerjee et al., 2019) or using shorter survey modules which can reveal network structure probabilistically (Breza et al., 2020). This paper demonstrates how data from commonly used digital technologies can reveal rich information about social network structure in a developing country context. Our targeting results come from simulations of a rich empirical model, where part of the network has already adopted. Nodes benefit incrementally from each of their contacts' adoption, by an amount that depends on their eventual communication. This network demand system accounts for the net social benefits of each node's adoption, so that maximizing diffusion may not be optimal if consumers do not internalize the social benefits or costs of adoption.

2. Background

We consider Rwanda during the period 2005–2009, during which the mobile network was expanding. The regulator restricted entry, so the mobile operator whose data we use held above 88% of the market, and its records reveal nearly the entirety of the country's remote communication.[3] During this period, almost all phones were basic phones used primarily for calling; mobile internet and mobile money were not available. Households that owned phones were significantly wealthier.

2008 adoption subsidy program

The Rwanda Utilities Regulatory Authority collects 2% of all operator revenues into a Universal Access Fund, to be used for programs that accelerate Information and Communication Technologies (ICTs). Many of these programs were targeted at rural connectivity. In one of these programs, the Rwandan government in 2008 purchased 53,352 handsets (amounting to roughly 8% of the country's stock of handsets at the time) and distributed them to individuals through local governments at a reduced price. The handsets were all the same model produced by Motorola, which was chosen because it was low cost and had a long battery life. The government was able to lower the price of the handsets by buying bulk at wholesale prices, and also offered a repayment plan.

Fifteen of 30 districts participated in the program. As shown in Table 1, participating districts tend to be poorer and rural, with low baseline mobile phone adoption (in participating districts, 4% of households had mobile phones in 2005, versus 12% in nonparticipating districts). Districts that were allocated handsets are shaded in Fig. 1a.

Handsets were allocated using decentralized distribution with self-targeting. Each district handled its own distribution; generally, individuals came to the district office to add their name to a list, districts requested that number of handsets, and then individuals received a handset a few months later. Participating districts received enough for between 1% and 15% of households.

Beneficiaries were to pay a fraction of the value of the handset (retail value approximately $24) through monthly repayments, but few of these payments were made. The bank financing the program was not able to provide details on the amount of payments made, and thus the net subsidy that each recipient ultimately received. We will consider results under three alternate assumptions about the net amount of the subsidy. Under a low subsidy amount, we assume that recipients made all payments and received only a wholesale discount of $7.[4] Under a high subsidy amount, recipients made no payments on a plan specified in New Times (2008), and thus received a $17 discount. And under a moderate amount, an amount halfway between, of $12. Based on the comments from implementers our primary results focus on this moderate value, and we report results under the extreme assumptions in the Appendix.

After receiving a subsidized handset, the recipient paid standard prices for calls on a prepaid plan (incoming calls were not charged).

3. Data

This project uses several data sources[5]:

Call detail records (CDR): We use anonymous call records from the dominant Rwandan operator, capturing nearly every call made over 4.5 years by the operator's mobile phone subscribers, growing from approximately 400,000 in January 2005 and to 1.5 million in May 2009. The data contains a list of transactions that can be represented as tuples: $(t, h, i, j, l_i, l_j, d)$, where t is the timestamp, h is the handset identifier, i is the account placing the call (an anonymized identifier corresponding to a phone number), j is the account receiving the call, l_x is the location of the tower used to transmit $x's$ end of the call, and d is the duration. We use this data to identify subsidy recipients, measure how they are connected to each other and to the rest of the network, and measure their usage of the subsidized handset.[6]

[3] There were few alternatives for remote communication: the fixed line network was small (with penetration below 0.4%), and mail service was insignificant.

[4] This amount obtained by comparing to the internal price of a handset available at retail with comparable features, the Motorola C117.

[5] See Björkegren (2019) for more details.

[6] May 2005, February 2009, and part of March 2009 are missing from this data.

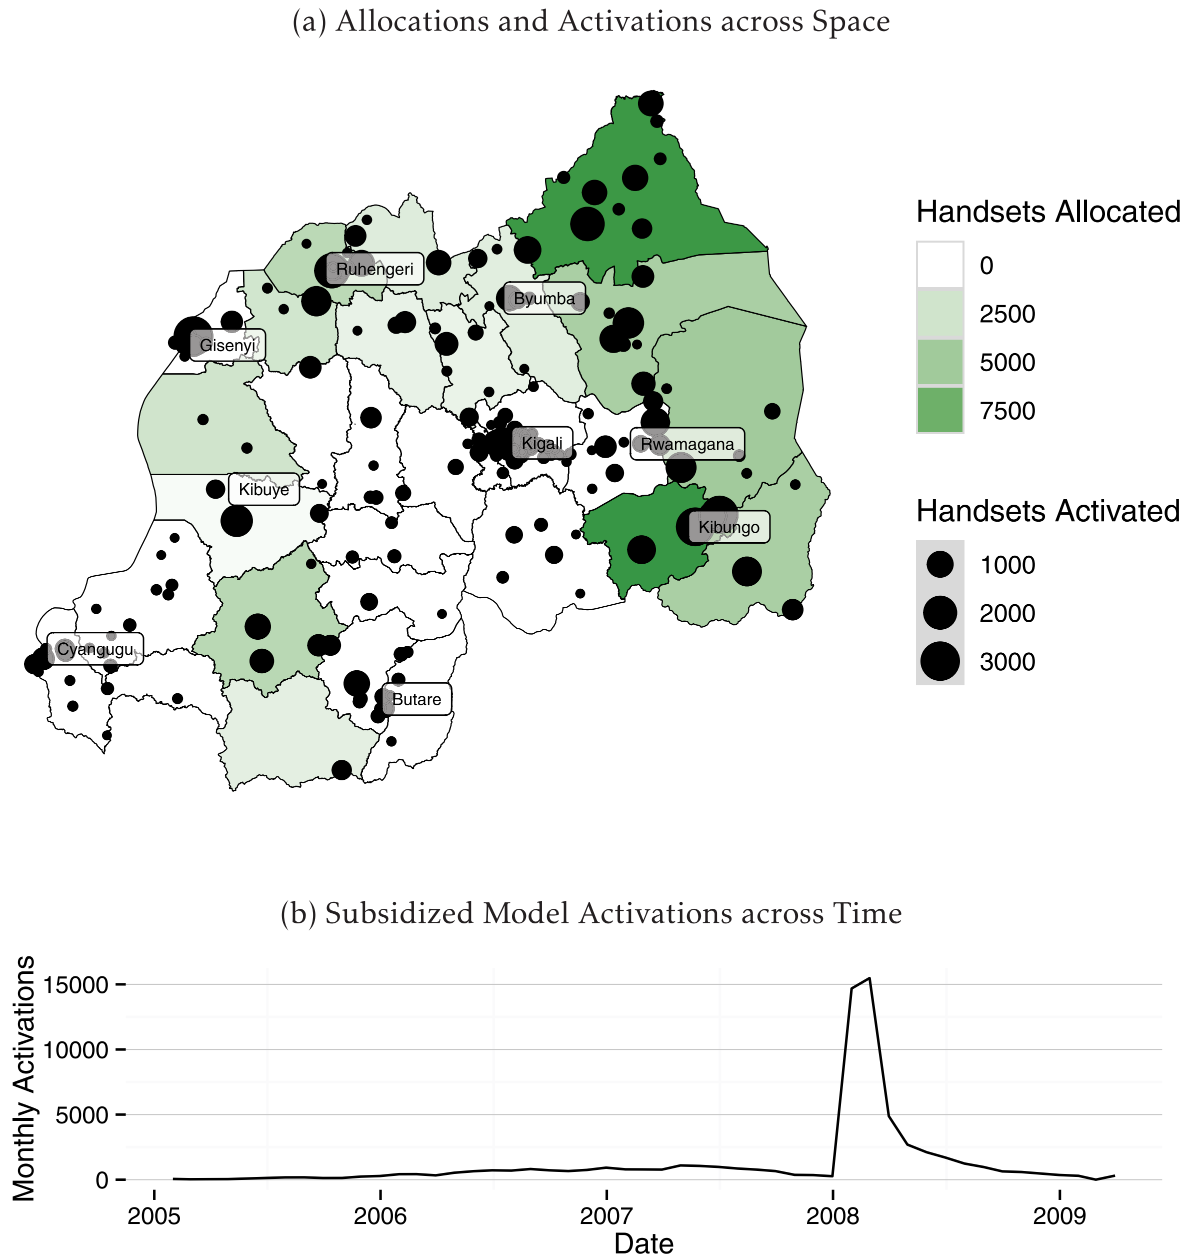

Source: Banque Rwandaise de Developpement. Activations from transaction data. Points represent the locations of cell phone towers where subsidized handsets were first used (locations are jittered). Major cities with large numbers of activations are labeled.

Fig. 1. Handset Subsidy Program.

Coverage: We create a baseline coverage map by computing the areas within line of sight of the towers operational in each month, a method suggested by the operator's network engineer. Elevation maps are derived from satellite imagery recorded by NASA (Jarvis et al., 2008; Farr et al., 2007).

Handset prices: We create a quality adjusted handset price index based on 160 popular models in Rwanda.

Household surveys: To compare districts that receive subsidies to those that do not, we use the National Institute of Statistics Rwanda's representative EICV surveys, fielded 2005–2006 (pre period) and 2010–2011 (post).

4. Descriptive evidence

4.1. Aggregate results

As a first step, one could analyze the impact of the program using national household survey data collected by the government in 2005 and 2010. We aggregate to the district level. A regression of the change in number of households owning phones on the number of handsets allocated provides correlational evidence of the impact of the

Table 2
Subsidy allocation and change in phone ownership.

	Proportion of households owning phones	
Post * Proportion allocated handsets in 2008	1.88 (0.95)	3.31 (1.06)
Post	0.35 (0.01)	0.32 (0.02)
District FEs	×	×
Districts included:	All	Participating
Mean in 2005	0.08	0.04
R^2	0.98	0.97
N districts	30	15
N	60	30

Proportion of households owning phones by district in 2005 and 2010, per Rwanda's EICV survey. Post equals 1 in 2010 and 0 in 2005. Proportion allocated handsets is total allocation divided by district population in 2002 census. Estimates computed using ordinary least squares, weighted by district population in 2002. Robust standard errors in parentheses.

program.[7] For a good with no network or learning effects, on average for every good that was allocated, we would expect to find an increase of less than one in a follow up survey, because the good may have deteriorated, been transferred, or simply offset a planned purchase (if it was inframarginal). If it spurs network effects, the allocation of one good may cause others to adopt, and could result in more than one good owned in a follow up survey.

Results are presented in Table 2. Being allocated handsets for an additional percentage point of households in a district in 2008 is associated with an increase of adoption between 1.88 and 3.31 percentage points between 2005 and 2010. Districts allocated handsets had lower initial levels of adoption, and there is evidence that adoption grew more slowly in areas with lower initial adoption: restricting to districts that received some handsets from the program tends to raise estimates.

However, allocations were not random and may have been targeted towards districts that otherwise would have differential adoption trends. Additionally, network effects need not remain constrained within a district; spillovers across district borders would bias estimates downward. Traditional datasets are coarse, making it difficult to investigate these assumptions or impacts.

The following sections evaluate the impact of the program using digital data. Our results suggest that there were indeed spillovers, and some of these spillovers benefited urban locations.

4.2. How were handsets used?

We first show how digital data can assess targeting of the 2008 Rwandan rural handset subsidy.

We infer which handsets are subsidized by the model, defining $subs(h) \in \{0, 1\}$. The particular model (Motorola C113) was otherwise rare in the country at the time, so we are able to identify beneficiaries based on receiving this model of handset during the dates of distribution. Fig. 1b shows activations of this model over time, showing a spike at the time of the subsidy. More than the allocated number of this model were eventually activated, suggesting some were eventually obtained outside the subsidy program. To be conservative, we consider an account as subsidized if its mode handset was the subsidized model, and it was activated during the first four months of 2008.[8] This gives us 41,225 accounts, 77% of the proposed allocation. Our simulations will report the impact of the program on this subset of the proposed allocation, and conservatively assume that the impact of the remaining handsets that are difficult to identify was zero.[9]

4.2.1. Stylized facts

Most handsets remained in rural areas, though there was some leakage to urban areas. 81% of the lines associated with subsidized handsets were activated in rural areas. Fig. 1a plots where handsets were allocated based on government records, as well as the number of subsidized handsets activated at each tower. Table 3 in row 2 shows that subsidy recipients (column 1) are much more likely to have their first tower in a rural area than the average of accounts that adopted in the same months (50%, column 2), or all phone accounts (42%, column 3). Subsidized handsets are also mostly used in rural areas. Because handsets are mobile, an individual may make calls from several locations, such as a village and the capital. The most used (modal) tower used for calls is rural for 85% of subsidy recipients (vs. 63% overall in those adoption months, and 55% for all accounts), as shown in row 3 of Table 3. This is suggestive of a program that for the most part reached people in intended rural areas.[10]

Subsidized handsets that were activated were used, not wasted. A common concern with subsidies is that goods may be allocated to consumers who do not value them. Because every transaction on handsets can be observed, we can assess this directly. Subsidy recipients use handsets in a similar manner as those who purchased phones around the same time. Table 3 shows that the ultimate recipients of subsidized handsets (column 1) use them on par with individuals who paid retail for phones around the same time (column 2), in terms of calls, durations, and total number of contacts. Usage is slightly less than the average over all accounts (column 3). (This is similar to Cohen and Dupas (2010)'s finding that subsidizing bednets did not affect the extent to which they are used.) This suggests that subsidized handsets ultimately were used by relatively typical users. While these users appear to have valued them, it is not clear if they would have also adopted in absence of the subsidy.

The implemented subsidy reached individuals with similar social networks as retail adopters. The handsets may have also induced

[7] The earlier survey does not ask about the number of handsets owned within a household, so we only look at the fraction of households owning at least one handset. If the subsidized handsets were distributed to households with existing handsets, this would underestimate the association.

[8] We define activation as the first transaction transmitted by the line on the phone network.

[9] We consider both the case where the remaining handsets have zero net benefit, and the case where they are wasted, so the government still incurs the cost, but the economy obtains no benefit.

[10] We do find that a small number of handsets passed through middlemen. The data include signatures of resale. One may see a particular handset h activated by a phone number i but subsequently passed along to phone number i'. Or, i' may be a middleman who briefly uses the handset for testing, and then transfers it to an ultimate user i''. We define a middleman as an account i' that uses two or more subsidized handsets h for 20 or fewer transactions in between two other accounts. We find 624 subsidized handsets (1.5% of those we identified) were transferred through 291 middlemen. We do not explicitly model these transfers.

Table 3
How phone accounts are used.

		Subsidy recipients Adopting 1-5.2008	Accounts Adopting 1-5.2008	Accounts All
Number		41,225	309,379	1,503,369
Rural	First tower, Mean	0.81	0.50	0.42
	Mode tower, Mean	0.85	0.63	0.55
Calls per month	Mean	37.7	37.5	40.0
	Median	28.7	26.1	24.1
	SD	34.0	48.9	59.0
Duration minutes per month	Mean	16.4	18.1	27.6
	Fraction to accounts subscribing after 1.2008	35%	33%	24%
	SD	23.0	47.1	92.2
Number of contacts (Degree)	Mean	62.2	57.5	105.8
	SD	42.8	73.4	159.9
Clustering coefficient	Mean	0.082	0.081	0.068
	SD	0.057	0.070	0.066

A location is considered rural if it is further than 15 km from the top 10 cities in the country.

spillovers, which would be mediated by the social network. We analyze the social network structure of recipients, as revealed by their later phone calls. We observe the communication graph G^T, where a directed link $ij \in G^T$ indicates that phone number i has called j by the last period of data, T (May 2009). This network includes only eventual phone subscribers, and so tends to be a wealthier segment of the population.

Recipients are similar to retail adopters in terms of network structure. As shown in Table 3, the fraction of a node's neighbors who are themselves connected (clustering coefficient) is very similar: 0.082 for subsidy recipients and 0.081 for all subscribing in the same months. One measure of how much recipients might tip the adoption of others is the eventual duration spoken with contacts that have yet to subscribe—these are people who might potentially adopt in response. This measure is very similar for recipients and nonrecipients (35% for subsidy recipients, 33% for all subscribing in the same months).

Recipients are also similar to retail adopters in terms of edges. Table 4 shows the properties of these edges. The first column presents the properties of edges in the entire network of subscribers, the second column shows edges from recipients to any subscriber, and the final column shows links between subsidy recipients. On average, 0.26 calls are placed per month, with a total duration that averages to 7.63 s. The average edge connects subscribers who are 30 km apart. 28% of edges only have calls during working hours. Most calls are short: 78% of edges have only had calls under 1 min, and 51% have only had calls under 30 s. The next two columns restrict these measures to subsidy recipients: first, edges from recipients to any subscriber, and second, only edges that connect two recipients. Edges connect subsidy recipients at shorter distances (average of 20 km) but otherwise these relationships appear similar to general calling relationships.

Altogether, these suggest that the implemented allocation (decentralized distribution with self-targeting) reached individuals with relatively typical social networks, not necessarily those with network properties particularly suited to spillovers.

The implemented subsidy reached individuals who are relatively connected to each other. Among all eventual subscribers, 2% of links are with subsidy recipients. Subsidy recipients themselves are nearly 5 times more connected to each other: 9% of the links of subsidy recipients are to other subsidy recipients. This suggests that decentralized distribution with self-targeting reached people who were relatively clustered in the social network.

To understand the overall effects of the subsidy we next use a structural model.

5. Structural model

The ultimate impact on network adoption and welfare depends on the interaction of the recipients' adoption decision with the network

Table 4
How phone accounts are linked: Edge properties.

Subset of nodes:	All	Subsidy recipients	
Edges:	All	All	Within
Calls per month	0.26	0.32	0.36
Duration per month (seconds)	7.63	7.50	8.02
Distance (km)	30.13	29.56	20.02
Any calls during			
...workday	0.59	0.60	0.58
...weekend	0.49	0.50	0.50
...late night	0.08	0.07	0.08
...holidays	0.28	0.22	0.21
Only calls during working hours	0.28	0.27	0.24
Call length			
...all under 30 sec	0.51	0.57	0.55
...all under 1 min	0.78	0.84	0.80
Nodes	1,503,675	41,225	41,225
Edges	195.6 m	4.4 m	0.4 m

Average of given attribute for the edges in the given subgraph, from January 2008 onwards.

of benefit flows. Björkegren (2019) estimates a network demand system for the Rwandan mobile phone network, using the same data. We use this structural model to simulate how the enacted subsidy program affected equilibrium adoption, and evaluate alternate subsidy programs. The model accounts for how a change to one individual's adoption affects others, which recursively affects others until reaching a new equilibrium. Our model does not explicitly account for resale; it assumes that any benefits accrue to the ultimate recipient of the phone.

5.1. Model

We briefly describe the empirical model of handset adoption that we use from Björkegren (2019). For more details, see that paper. The utility of owning a phone is derived from making calls, so we begin with a model of usage.

5.1.1. Usage

Let S_t be the subset of nodes subscribing in month t. At each period t, individual i can call any contact j that currently subscribes, $j \in G_i^T \cap S_t$, to receive utility u_{ijt}. Each month, i draws a communication shock $\epsilon_{ijt} \overset{iid}{\sim} F_{ij}$ representing a desire to call contact j. Given the shock, i chooses a total duration $d_{ijt} \geq 0$ for that month, earning utility:

$$u_{ijt} = \max_{d_{ijt}\geq 0}\left[\frac{1}{\beta_{cost}} v(d_{ijt}, \epsilon_{ijt}) - c_{ijt} d\right]$$

The benefit of making calls is:

$$v(d,\epsilon) = d - \frac{1}{\epsilon}\left[\frac{d^\gamma}{\gamma} + \alpha d\right]$$

where the first term represents a linear benefit; $\gamma > 1$ controls how quickly marginal returns decline, and α controls the intercept of marginal utility, and thus the fraction of months for which no call is placed.

The marginal cost is:

$$c_{ijt} = p_t + \beta_{coverage}\phi_{it}\phi_{jt}$$

where p_t is the per-second calling price (including any tax), and $\beta_{coverage}\phi_{it}\phi_{jt}$ represents a hassle cost when the caller or receiver have imperfect coverage. An individual's coverage $\phi_{it} \in [0,1]$ is derived from the fraction of the area surrounding his most used locations receiving cellular coverage in month t.

The expected utility i receives from being able to call j in time period t is given by:

$$\mathbb{E}u_{ij}(p_t,\boldsymbol{\phi}_t) = \int_{\underline{\epsilon}_{ijt}}^{\infty}\left[d(\epsilon,p_t,\boldsymbol{\phi}_t)\cdot\left(\frac{1}{\beta_{cost}}\left(1-\frac{\alpha}{\epsilon}\right) - p_t - \beta_{coverage}\phi_{it}\phi_{jt}\right) - \frac{1}{\beta_{cost}\epsilon}\frac{d(\epsilon,p_t,\boldsymbol{\phi}_t)^\gamma}{\gamma}\right]dF_{ij}(\epsilon)$$

where $\boldsymbol{\phi}_t$ represents the vector of coverage for all individuals, and the smallest shock at which i would call j is given by $\underline{\epsilon}_{ijt} := \frac{1+\alpha}{1-\beta_{cost}\left(p_t+\beta_{coverage}\phi_{it}\phi_{jt}\right)}$.[11]

Altogether, each month i is on the network, he receives expected utility from each contact who is also on the network:

$$\mathbb{E}u_{it}(p_t,\boldsymbol{\phi}_t,\boldsymbol{x}_{Gi}) = \sum_{j\in G_i \text{ and } x_j\le t}\mathbb{E}u_{ij}(p_t,\boldsymbol{\phi}_t)$$

where x_j represents j's adoption time. Each month that i is not on the network he receives utility zero.

5.1.2. Adoption

At period t, i knows the current price of a handset, $p_{it}^{handset}$ (including any tax).[12] He believes his contacts will adopt at times $\boldsymbol{x}_{G_i}$, and that in period $x > t$, the handset price will be $\mathbb{E}_t p_{ix}^{handset}$, a deterministic function that is described in the following section. He expects the utility of adopting at time x to be:

$$\mathbb{E}_t U_i^x(\boldsymbol{x}_{G_i}) = \delta^x\left[\sum_{s\ge x}^{\infty}\delta^{s-x}\mathbb{E}u_{is}(p_s,\boldsymbol{\phi}_s,\boldsymbol{x}_{G_i}) - \mathbb{E}_t p_{ix}^{handset} + \eta_i\right]$$

i adopts at the first month x_i where he expects adopting immediately to be more attractive than waiting:

$$\min\ x_i\ s.t.\left[\mathbb{E}_{x_i}U_i^{x_i}(\boldsymbol{x}_{G_i}) \ge \max_{s>x_i}\mathbb{E}_{x_i}U_i^s(\boldsymbol{x}_{G_i})\right] \tag{1}$$

An individual's type η_i captures any residual heterogeneity explaining why i adopted at x_i.

5.1.3. Network adoption equilibrium

Initial adopters (S_0) are held fixed. Each other individual i decides on an adoption time $x_i \in \{1,\dots,\bar{T}\}$, for some $\bar{T} \ge T$.

An **equilibrium** Γ is defined by adoption dates $\boldsymbol{x} = \left[x_i\right]_{i\in S}$ such that each individual $i \in S\backslash S_0$ adopts optimally according to Eq. (1), with beliefs consistent with when his contacts adopt.

Individuals correctly forecast call prices p_x, coverage $\boldsymbol{\phi}_x$, and the dates their contacts adopt $\boldsymbol{x}_{G_i}$. Because a handset becomes sunk at the time of purchase, forecasts of future handset prices can sway the adoption decision. We assume that at each period t, individuals learn their current handset price, which can be decomposed into the market price minus any potential subsidy if they are targeted:

$$p_{it}^{handset} = p_t^{handset} - \Delta\cdot \mathrm{I}(i \text{ subsidized at time } t)$$

Δ represents the net present discount resulting from the subsidy payment plan; the spacing of real payments over time does not influence adoption. They expect the base price in future periods to decline at an exponential rate consistent with the overall decline over this period. When adopting, they anticipate any subsidy they will receive:

$$\mathbb{E}_t p_{ix}^{handset} = \omega^{x-t}p_t^{handset} - \Delta\cdot \mathrm{I}(i \text{ subsidized at time } x)$$

for $\omega = \left(\frac{p_{\bar{T}}^{handset}}{p_0^{handset}}\right)^{\frac{1}{\bar{T}}}$.

Under this model, each individual's adoption and usage depends on the adoption decisions of his contacts, which in turn depend on the adoption decisions of her contacts, and so on. A perturbation of utility that causes one individual to change their adoption date can shift the equilibrium, inducing ripple effects through potentially the entire network.

5.1.4. Firm and government

The firm earns revenue from the price it charges for calls (p_t), and the government earns revenue from taxes on adoption ($\tau_{it}^{handset}$) and usage (τ_{it}^{usage}). The paths of these rates are held fixed and announced in advance.[13]

5.2. Estimation

The parameters of the model are estimated from data in Björkegren (2019). Individuals choose when to adopt a mobile phone and, if they adopt, how to use the phone. The decision to use a phone directly reveals the value of each connection, overcoming traditional issues with identifying the value of network goods solely from the decision to adopt. Specifically, i's decision to call j for $d_{ijt} \ge 0$ seconds in month t when facing prices p_t and coverage ϕ_{it} and ϕ_{jt} identifies sensitivity to costs (monetary β_{cost} and the hassle of imperfect coverage $\beta_{coverage}$),

[11] This model assumes that the utility of a call accrues to the person who pays for it (the caller), rather than model how call utility is split between caller and receiver. Björkegren (2019) also considers the possibility that consumers additionally earn the equivalent utility from the calls they receive (which are free); although market outcomes are not very different for the counterfactuals in that article, this double counts call utility relative to the utility implied by the adoption decision.

[12] Note the handset price has an i subscript as it may include individual specific subsidies, different from Björkegren (2019).

[13] The government earns

$$R_G^\Gamma = \sum_{i\in S \text{ and } x_i\le T}\left[-\Delta\cdot\delta^{x_i}\cdot 1_{\{i\text{subsidized and takes up subsidy}\}} + \delta^{x_i}\tau_{ix_i}^{handset}p_{ix_i}^{handset} + \sum_{t\ge x_i}^{T}\delta^t\tau_{it}^{usage}p_t\sum_{j\in G_i^T\cap S_t}\mathbb{E}d_{ij}(p_t,\boldsymbol{\phi}_t)\right]$$

and the firm earns revenue

$$R_F^\Gamma = \sum_{i\in S \text{ and } x_i\le T}\sum_{t\ge x_i}^{T}\delta^t(1-\tau_{it}^{usage})p_t\sum_{j\in G_i^T\cap S_t}\mathbb{E}d_{ij}(p_t,\boldsymbol{\phi}_t)$$

where the expected duration of calls from i to j is given by

$$\mathbb{E}d_{ij}(p_t,\boldsymbol{\phi}_t) = \int_{\underline{\epsilon}_{ijt}}^{\infty}d(\epsilon,p_t,\boldsymbol{\phi}_t)\cdot dF_{ij}(\epsilon)$$

the call shock distributions (F_{ij}), and parameters affecting the shape of the utility function (γ and α).[14]

The adoption decision identifies bounds on idiosyncratic preferences for having a phone, $[\underline{\eta}_i, \bar{\eta}_i]$. At time x_i, i bought a handset rather than waiting K months. This implies that the expected utility of being on the network during the following K months must have exceeded the expected drop in handset prices. Similarly, i could have purchased a handset earlier. At time $x_i - K$, i chose to wait, so he must have preferred some future adoption date. Those months provided less expected utility than the expected drop in handset prices.

These inequalities imply bounds for each individual's type, $\underline{\eta}_i \leq \eta_i \leq \bar{\eta}_i$. We compute these as follows:

$$\begin{aligned}\underline{\eta}_i = & \frac{1}{1-\delta^K}\Bigg[p^{handset}_{i,x_i} - \delta^K \mathbb{E}_{x_i} p^{handset}_{i,x_i+K} \\ & - \sum_{s=0}^{K-1} \delta^s \mathbb{E} u_{i,x_i+s}(p_{x_i+s}, \boldsymbol{\phi}_{x_i+s}, \boldsymbol{x}_{G_i})\Bigg] \\ \bar{\eta}_i = & \max_{\tilde{K}>0}\Bigg[\frac{1}{1-\delta^{\tilde{K}}}\Bigg[p^{handset}_{i,x_i-K} - \delta^{\tilde{K}} \mathbb{E}_{x_i-K} p^{handset}_{x_i-K+\tilde{K}} \\ & - \sum_{s=0}^{\tilde{K}-1} \delta^s \mathbb{E} u_{i,x_i-K+s}(p_{x_i-K+s}, \boldsymbol{\phi}_{x_i-K+s}, \boldsymbol{x}_{G_i})\Bigg]\Bigg]\end{aligned} \tag{2}$$

These affect the adoption choice, but we do not include them in welfare computations because they may be a function of beliefs.

Recovering these bounds is straightforward for accounts that purchased phones at retail prices. However, it is more nuanced for recipients of the enacted subsidy program. We make the following assumptions:

1. In the baseline, all eligible individuals took up the subsidy[15]
2. Recipients preferred taking the subsidy at the point of adoption to purchasing any time in the following 4 years[16]
3. In our main results we assume that the present discounted value of the subsidy is \$12.00, as described in Section 2. In the Appendix we show results under the assumption that the subsidy was \$6.81 or \$17.19.
4. Recipients did not delay adoption in anticipation of the subsidy[17]

The standard upper bound $\bar{\eta}_i$ would admit the possibility that subsidy recipients valued phones so much that they delayed adoption in order to obtain the subsidy—that would imply the program had a negative effect on adoption. Conversations with designers and implementers of the program suggested it is unlikely that recipients valued phones that much. The last assumption above implies that we instead set $\bar{\eta}_i$ to the level had i not delayed adoption in anticipation of the subsidy.

This estimate of $\bar{\eta}_i$ still admits high valuations for phones among recipients: it mechanically suggests that recipients would have adopted at the exact same month regardless of the subsidy.[18] Because that mechanical bound is not very informative, and because subsidy recipients are observably similar to nonrecipients, we consider the bound $\bar{\eta}_i^{focal}$, which assumes that subsidy recipients' unobserved value for having a handset is on average no higher than those who paid for phones in the same months. For nonrecipients, the focal type estimate $\bar{\eta}_i^{focal}$ is set to the upper bound $\bar{\eta}_i$; for recipients it is set to a weighted average of the upper and lower bound: $\bar{\eta}_i^{focal} = a\underline{\eta}_i + (1-a)\bar{\eta}_i$. We set the weight a so that the average of $\bar{\eta}_i^{focal}$ coincides with the average of $\bar{\eta}_i$ for nonrecipients who subscribed during those months:

$$a = \frac{\sum_{i \in S_{unsubs0801-0805}} \bar{\eta}_i / N_{unsubs0801-0805} - \sum_{i\in S_{subs}} \bar{\eta}_i / N_{subs}}{\sum_{i \in S_{subs}} \underline{\eta}_i / N_{subs} - \sum_{i \in S_{subs}} \bar{\eta}_i / N_{subs}}$$

This nuance has less of an effect on the evaluation of counterfactual targeting regimes, since types are more precisely revealed for nonrecipients' later purchase decisions.

5.3. Simulation

We use the iterated best response method developed in Björkegren (2019) to first simulate adoption and usage in the baseline environment, and then under counterfactuals. The method is initialized with a candidate adoption path $\boldsymbol{x}^0$, from which each individual sequentially reoptimizes their adoption date, conditional on the adoption dates of others, until the path converges.

There are two dimensions of uncertainty. As discussed above, types are set identified rather than point identified. Additionally, there tend to be multiple equilibria because individuals may coordinate on being optimistic or pessimistic about others' adoption. This equilibrium selection is controlled by the candidate adoption path.

As in Björkegren (2019), we consider outcomes Y at the upper bound of the highest equilibrium ($Y^{\bar{\Gamma}}$ with types $\bar{\eta}$, initialized with non initial adopters adopting immediately) and lower bound of the lowest equilibrium ($Y^{\underline{\Gamma}}$ and $\underline{\eta}$, initialized with non initial adopters delaying adoption to the last period).[19] We additionally select a *focal* equilibrium, using the type estimates $\bar{\eta}_i^{focal}$ and initializing with the adoption path that was observed in the data. This will tend to select an equilibrium close to the one that was observed. We measure policy impacts by reporting, for each specified equilibrium, the difference between that equilibrium with and without the subsidies.[20]

6. Impact of actual subsidy program

In Table 5 we compute the baseline simulation with decentralized distribution with self-targeting ("with subsidy" in the table), as well as simulations where the subsidy has been removed. We first allow each recipient to reoptimize their decision individually, without allowing those changes to ripple through the network ("proximal effect"). We then allow all nodes to adjust their decisions until a new equilibrium is reached ("additional ripple effect"). The combined effect is reported

[14] The distribution for call shocks is parameterized as a mixture distribution: $F_{ij}[\epsilon] = q_i \Phi\left(\frac{\ln(\epsilon)-\mu_{ij}}{\sigma_i}\right) + (1-q_i)1_{\{\epsilon > -\infty\}}$, where $\Phi(\cdot)$ represents the standard normal CDF. The first component is a lognormal distribution, $\ln N(\mu_{ij}, \sigma_i^2)$, and the second a point mass, under which there are no calls regardless of the cost.

[15] Given the decentralized nature of the implemented subsidy program, it is difficult to determine the entire set of individuals who were eligible. Since the subsidy was very attractive, we assume that all eligible individuals took up the subsidy and that it was valid only in the month they adopted.

[16] That is, for recipients we compute $\underline{\eta}_i = \min_{K \in \{1,\ldots,48\}} \underline{\eta}_i(K)$.

[17] Additionally, for subsidy recipients, we back out bounds on types under the simulated baseline adoption paths of nonrecipients, rather than the observed adoption path in the data, as the latter would lead to some recipients not taking up the subsidy due to noise.

[18] This is because it is unaffected by whether i is subsidized.

[19] Because there is a monotonic relationship between adoption date and call utility, these outcomes bound all possible equilibrium outcomes for call utility and firm revenue R_F. The high and low equilibrium outcomes of U_{net} and R_G may not bound all possible equilibrium outcomes, though deviations appear to be minor. Because the net utility function omits idiosyncratic benefits, it does not match the utility each individual maximizes; there may be an equilibrium between $\underline{\Gamma}$ and $\bar{\Gamma}$ that has a net utility lying outside the bounds of $U_{net}^{\underline{\Gamma}}$ and $U_{net}^{\bar{\Gamma}}$. Similarly because handset prices are decreasing, government revenue may be a nonmonotonic function of adoption date, and there may be an equilibrium between $\underline{\Gamma}$ and $\bar{\Gamma}$ that generates government revenue outside the bounds of $R_G^{\underline{\Gamma}}$ and $R_G^{\bar{\Gamma}}$.

[20] A more natural measure of policy impact would be bounds on the changes in revenue and utility across the range of equilibria; however, this measure is computationally prohibitive because adoption decisions are interlinked. Björkegren (2019), in the Supplemental Appendix, finds that the changes in the upper and lower bounds often bound the true impact in Monte Carlo simulations.

Table 5
Baseline and impact of implemented adoption subsidy program.

	All nodes			Recipients			Nonrecipients		
Number	1,503,675			41,225			1,462,450		
	Lower	Focal	Upper	Lower	Focal	Upper	Lower	Focal	Upper
Adoption time (mean, month)									
Baseline with subsidy	[24.10,	<u>22.11</u>,	22.10]	[37.38,	<u>37.38</u>,	37.38]	[23.73,	<u>21.68</u>,	21.67]
Total impact of subsidy	−0.60	<u>−0.04</u>	−0.00	−18.01	<u>−1.41</u>	−0.00	−0.11	<u>−0.01</u>	−0.00
... proximal effect	−0.46	<u>−0.04</u>	−0.00	−16.94	<u>−1.29</u>	−0.00	−0.00	<u>−0.00</u>	−0.00
... additional ripple effect	−0.13	<u>−0.01</u>	−0.00	−1.07	<u>−0.11</u>	−0.00	−0.11	<u>−0.01</u>	−0.00
Consumer surplus (total, million $)									
Baseline with subsidy	[243.51,	<u>269.59</u>,	269.70]	[1.82,	<u>1.88</u>,	1.88]	[241.68,	<u>267.71</u>,	267.82]
Total impact of subsidy	3.68	<u>0.60</u>	0.40	1.03	<u>0.44</u>	0.40	2.65	<u>0.17</u>	−0.00
... proximal effect	2.54	<u>0.55</u>	0.40	0.95	<u>0.43</u>	0.40	1.59	<u>0.12</u>	−0.00
... additional ripple effect	1.13	<u>0.05</u>	−0.00	0.08	<u>0.01</u>	−0.00	1.06	<u>0.04</u>	−0.00
Firm revenue (total, million $)									
Baseline with subsidy	[165.07,	<u>187.26</u>,	187.37]	[0.88,	<u>0.91</u>,	0.91]	[164.19,	<u>186.35</u>,	186.46]
Total impact of subsidy	1.82	<u>0.15</u>	−0.00	0.55	<u>0.05</u>	−0.00	1.27	<u>0.10</u>	−0.00
... proximal effect	1.23	<u>0.11</u>	−0.00	0.51	<u>0.04</u>	−0.00	0.72	<u>0.07</u>	−0.00
... additional ripple effect	0.59	<u>0.04</u>	−0.00	0.04	<u>0.01</u>	−0.00	0.54	<u>0.03</u>	−0.00
Government revenue (total, million $)									
Baseline with subsidy	[65.36,	<u>73.09</u>,	73.13]	[0.38,	<u>0.39</u>,	0.39]	[64.98,	<u>72.70</u>,	72.74]
Total impact of subsidy	0.50	<u>−0.34</u>	−0.40	0.09	<u>−0.37</u>	−0.40	0.41	<u>0.03</u>	−0.00
... proximal effect	0.27	<u>−0.35</u>	−0.40	0.05	<u>−0.37</u>	−0.40	0.22	<u>0.02</u>	−0.00
... additional ripple effect	0.23	<u>0.01</u>	−0.00	0.04	<u>0.00</u>	−0.00	0.20	<u>0.01</u>	−0.00

Results in each cell reported for the lower, <u>focal</u>, and upper bound estimate of the equilibrium. Impacts represent the difference in these bounds. We hold fixed the adoption decision of 9 subsidized nodes that have crossed bounds for η_i. Utility and revenue reported in 2005 U.S. Dollars, discounted at a monthly rate of $\delta \approx 0.9945$. Consumer surplus includes the surplus utility each individual receives from the call model through May 2009, minus the cost of holding a handset from the time of adoption until May 2009.

under "Total Impact of the Subsidy". The first column shows the results for all nodes; subsequent columns show results for subsidy recipients and nonrecipients.

Because the decision to purchase a subsidized good only loosely reveals how much the recipient values it, the bounds for each equilibrium are wide. The lower bound of the lower equilibrium presents a pessimistic scenario: targeted individuals would have delayed adoption by an average of 18 months, and the upper bound of the upper equilibrium admits the possibility of zero impact: it mechanically suggests that targeted individuals would not have delayed adoption at all in absence of the subsidy. The focal equilibrium presents a more reasonable scenario: targeted individuals would have delayed adoption by 1.4 months on average in the absence of the subsidy.

We find that in the focal equilibrium:

The subsidy improved welfare. Factoring in the net present cost of the subsidy of $0.46 m, it shifted net welfare upward $0.41 m, a social rate of return of 89%. That calculation assumes the program's benefits arise only from the 41,225 recipients we can clearly identify and the net benefit to the remaining 12,127 handsets was zero. Even if the value of these remaining handsets was completely destroyed through misallocation, the program would still have a social rate of return of 44%.

A substantial fraction of benefits accrued to *non*recipients. Recipients' utility increased by $0.44 m, from the combination of increased calling and the direct value of the discount. Nonrecipients only received utility from increased calling, but obtained 25% of all consumer surplus, due to spillovers.

Nonrecipients account for most of the increase in revenue. Nonrecipients account for 65% of the revenue from the subsidy.

Ripple effects were important. Ripple effects account for 24% of the effect on revenue and 8% of the effect on consumer surplus.

Net welfare effects are similar if we assume that the implemented subsidy was a different amount (see Appendix Table 7). Next, we consider alternative targeting rules, and break down impacts on urban and rural consumers.

7. Alternative targeting

Decentralized distribution with self-targeting is one of many potential targeting schemes that the government could have used. This section evaluates other schemes. We hold fixed the subsidy amount and number of nodes allocated (41,225), implement it in January 2008, and vary who receives the subsidy. All schemes allocate the subsidy to nodes that did not receive the actual subsidy, that had yet to adopt by January 2008, but that had adopted by the end of the data (May 2009).

Table 6 reports simulation results from alternate targeting rules. We report consumer surplus (rural and urban), firm revenue, net government revenue, and their sum (net welfare) in the lower, focal, and upper bound equilibria. We consider variants of three types of targeting rules. The first two are theoretical allocations that rely on network information that we observe, but which would not have been known at the time.

In some network problems it can be beneficial to disperse benefits throughout the network. Such dispersed interventions could cause marginal nodes to tip in their respective regions of the network. *Random* (rows 2–5) employs such a strategy, selecting nodes at random throughout the network of those who eventually adopted.[21]

There is also substantial interest in targeting individuals who are particularly influential, or well-connected. We consider two such strategies. *Degree* (rows 6) ranks nodes by degree (eventual connections), and then allocates the subsidy in order until the budget is depleted. This prioritizes people who have many immediate contacts, regardless of whether those contacts have many contacts. In simple diffusion models such people may induce more first order spillovers. *Centrality* (rows 7) ranks nodes by eigenvector centrality, a measure of connectedness to the entire network, and then allocates the subsidy in order until the budget is depleted. This will prioritize people who are connected to many connected people. In simple diffusion models such people may induce more higher order spillovers.

We also consider a practical variant of targeting rule that could be implemented. It is often difficult for central authorities to gather information about a network, and in our setting we only observe the network and usage after a person has adopted. The third strategy utilizes the fact that individuals tend to have information about

[21] We simulate from 3 random draws, and report the mean and standard deviation of each outcome across draws.

Table 6
Impact of alternate targeting policies.

Targeting policy		Welfare ($m)			Consumer surplus ($m)						Revenue ($m)					
		Net			Rural			Urban			Firm			Government		
		Lower	Focal	Upper	Lower	Focal	Upper	Lower	Focal	Upper	Lower	Focal	Upper	Lower	Focal	Upper
Implemented allocation																
Decentralized distribution with self-targeting		6.00	0.41	0.00	2.09	0.45	0.34	1.59	0.15	0.06	1.82	0.15	−0.0	0.50	−0.34	−0.40
Theoretical allocations																
Random, rural	Mean	2.64	0.21	0.35	1.11	0.49	0.53	0.66	0.02	0.05	0.91	0.07	0.13	−0.04	−0.38	−0.36
	SD	0.04	0.01	0.03	0.01	0.00	0.00	0.03	0.00	0.01	0.01	0.01	0.01	0.00	0.00	0.00
Random, urban	Mean	5.14	0.31	0.47	0.31	0.02	0.06	3.15	0.57	0.61	1.53	0.09	0.15	0.15	−0.37	−0.35
	SD	0.32	0.08	0.06	0.00	0.00	0.00	0.21	0.05	0.03	0.09	0.03	0.02	0.03	0.01	0.01
Highest degree		6.39	0.25	0.39	1.33	0.32	0.35	2.67	0.25	0.29	2.03	0.06	0.12	0.35	−0.38	−0.36
Highest eigenvector centrality		6.46	0.20	0.33	0.82	0.15	0.19	3.35	0.41	0.43	1.98	0.03	0.08	0.31	−0.39	−0.37
Practical allocations																
Vouchers		6.25	1.25	1.34	1.11	0.43	0.46	2.90	0.63	0.64	1.92	0.44	0.48	0.32	−0.25	−0.24
Vouchers, rural		5.82	1.18	1.27	1.29	0.52	0.54	2.43	0.50	0.51	1.81	0.41	0.45	0.29	−0.26	−0.24

Each cell reports the effect of subsidy policies, in the lower, focal, and upper equilibrium, for a subsidy amount of $12. Impacts represent the difference in these bounds. Implemented subsidy represents the negative of the impact of removing the subsidy from the baseline. For other policies the impact is of adding the counterfactual subsidy from the equilibrium where the implemented subsidy is removed. Utility and revenue reported in 2005 U.S. Dollars, discounted at a monthly rate of $\delta \approx 0.9945$. Consumer surplus includes the surplus utility each individual receives from the call model through May 2009, minus the cost of holding a handset from the time of adoption until May 2009. Rural accounts are defined as those whose modal tower is more than 15 km from a city.

the individuals around them. It exploits that private information to improve targeting. *Vouchers* (final rows) provides current subscribers with adoption vouchers for the subsidy amount that can be passed on. This is a variant of a nomination approach, which has been found to be successful in targeting health interventions (Kim et al., 2015). We assume that each selected subscriber gives this voucher to the contact he eventually talks the most with but who has yet to subscribe (his strongest unsubscribed link).[22] In the context of the model, this is typically the contact that the subscriber would earn the most utility from communicating with. This could be implemented in partnership with the phone company, by texting targeted subscribers a unique code that could be redeemed in the presence of this contact. Providing such vouchers is a common strategy used by tech companies that are growing user networks.

We find:

There are tradeoffs between theoretical allocations. The random and high degree allocations induce aggregate welfare effects in a similar range ($0.22–0.31 m in the focal equilibrium), suggesting that there may not be strong reasons to favor one over the other on those grounds alone. However, there is a trade-off between simple aggregate efficiency and reaching remote people. Allocations to random rural nodes (row 2) have aggregate welfare effects that are 32% lower than those to random urban nodes (row 4), in the focal equilibrium. Additionally, while targeting high degree nodes results in a fairly even split in benefits to rural and urban consumers (43% to urban consumers), targeting the most central nodes results in 73% of benefits going to urban consumers, who are more likely to be central. These results suggest that there may be efficiency/equity tradeoffs. Targeting remote nodes may require higher welfare weights on remote people.

Vouchers yield high welfare impacts. Vouchers outperform other allocations by aggregate measures (in the focal and high equilibrium, and obtain comparable performance to the best in the lowest), and achieve high performance on all components of welfare (last rows of Table 6). Relative to the baseline subsidy, in the focal equilibrium vouchers given to rural adopters improve net welfare by $0.77 m, and the surplus of both rural ($0.07 m) and urban ($0.35 m) consumers. In other contexts, the good performance of voucher allocations may result from the friendship paradox (Feld, 1991)—that the contacts of a randomly selected node will tend to be more central. However, here the high welfare impacts of vouchers arises not so much because they reach people who induce adoption spillovers (network structure), but rather because they reach people who will heavily use the phones (demand parameters). This can be seen in two ways. First, in the higher equilibria we find that the allocations that directly target centrality (degree and eigenvector) generate less social surplus than vouchers, suggesting that vouchers capture information different from those measures of centrality. Second, most of the effect of vouchers come from proximal effects on recipients rather than ripple effects on the rest of the network.[23] The prevalence of this type of marketing strategy in successful network startups may indicate that some combination of these features may be common in real world networks.

Targeted allocations induce spillovers across regions. Even when random subsidies are distributed only in rural areas, 4% of consumer surplus accrues to urban areas in the focal equilibrium; and 3% vice versa (rows 2 vs. 4). This results from both increased calls between the regions, and increased adoption in the spillover region.

Different network targeting rules work better in different equilibria. Network targeting rules (degree and centrality) perform much better than random in the lower bound of the lowest equilibrium, where consumers are pessimistic about others' adoption. Conversely, in higher equilibria, random allocations perform similarly and sometimes better than the network targeting rules.

Appendix Table 7 assesses the effects of subsidies of different amounts. Larger subsidies tend to yield higher aggregate welfare benefits, however even small vouchers yield high welfare gains, suggesting a high social rate of return.

7.1. Discussion

Decentralized distribution with self-targeting had an impact in the range of effective alternatives, and is not dominated across all equilibria by any of the theoretical allocations.

In counterfactual allocations, impacts are similar between the upper and focal equilibria. This highlights that subsidies do not necessarily have zero impact in the upper equilibrium—the zero effect for the

[22] We evaluate vouchers given to the earliest adopters (adopting January 2005 or prior); effects for recent adopters (adopting December 2007) are similar.

[23] By construction, vouchers target people who substantially benefit from using a phone (high eventual usage) but who have not yet adopted. The model attributes their non-adoption to idiosyncratic reasons, assigning a low η_i. Although η_i affects adoption, it is not included in welfare, as it may represent beliefs. The vouchers would induce lower welfare effects if in fact voucher recipients received less utility from having a phone despite using it heavily, or if subscribers gave vouchers to people other than their strong contacts.

Table 7
Impact of alternate targeting policies, alternate subsidy amounts.

Targeting policy		Subsidy	Welfare ($m)			Consumer surplus ($m)						Revenue ($m)					
		Amount ($)	Net			Rural			Urban			Firm			Government		
			Lower	Focal	Upper	Lower	Focal	Upper	Lower	Focal	Upper	Lower	Focal	Upper	Lower	Focal	Upper
Implemented subsidy		7	4.90	0.38	0.00	1.58	0.29	0.19	1.31	0.12	0.03	1.54	0.14	−0.0	0.47	−0.17	−0.23
		12	6.00	0.41	0.00	2.09	0.45	0.34	1.59	0.15	0.06	1.82	0.15	−0.0	0.50	−0.34	−0.40
		17	6.63	0.43	0.00	2.48	0.61	0.49	1.73	0.18	0.09	1.98	0.16	−0.0	0.45	−0.51	−0.57
Theoretical allocations																	
Random, rural	Mean	7	2.09	0.20	0.35	0.74	0.31	0.35	0.55	0.02	0.05	0.72	0.07	0.12	0.08	−0.20	−0.18
	SD	7	0.08	0.01	0.03	0.00	0.00	0.00	0.05	0.00	0.01	0.02	0.01	0.01	0.01	0.00	0.00
	Mean	12	2.64	0.21	0.35	1.11	0.49	0.53	0.66	0.02	0.05	0.91	0.07	0.13	−0.04	−0.38	−0.36
	SD	12	0.04	0.01	0.03	0.01	0.00	0.00	0.03	0.00	0.01	0.01	0.01	0.01	0.00	0.00	0.00
	Mean	17	2.84	0.21	0.36	1.37	0.67	0.70	0.69	0.02	0.05	0.97	0.07	0.13	−0.20	−0.55	−0.53
	SD	17	0.05	0.02	0.03	0.01	0.00	0.00	0.03	0.01	0.01	0.01	0.01	0.01	0.00	0.00	0.00
Random, urban	Mean	7	4.04	0.26	0.42	0.24	0.02	0.06	2.36	0.37	0.41	1.21	0.07	0.13	0.23	−0.20	−0.18
	SD	7	0.30	0.07	0.04	0.01	0.00	0.00	0.19	0.04	0.02	0.08	0.03	0.02	0.02	0.01	0.01
	Mean	12	5.14	0.31	0.47	0.31	0.02	0.06	3.15	0.57	0.61	1.53	0.09	0.15	0.15	−0.37	−0.35
	SD	12	0.32	0.08	0.06	0.00	0.00	0.00	0.21	0.05	0.03	0.09	0.03	0.02	0.03	0.01	0.01
	Mean	17	5.47	0.34	0.49	0.34	0.02	0.07	3.51	0.77	0.80	1.63	0.10	0.16	0.00	−0.54	−0.52
	SD	17	0.30	0.06	0.06	0.00	0.00	0.00	0.19	0.03	0.03	0.08	0.02	0.02	0.02	0.01	0.01
Highest degree		7	4.26	0.13	0.27	0.78	0.20	0.24	1.79	0.12	0.16	1.36	0.02	0.07	0.33	−0.22	−0.20
		12	6.39	0.25	0.39	1.33	0.32	0.35	2.67	0.25	0.29	2.03	0.06	0.12	0.35	−0.38	−0.36
		17	7.25	0.33	0.44	1.63	0.42	0.46	3.07	0.36	0.38	2.30	0.09	0.13	0.25	−0.55	−0.53
Highest eigenvector centrality		7	5.01	0.09	0.224	0.54	0.09	0.13	2.59	0.23	0.25	1.53	−0.01	0.0	0.35	−0.23	−0.21
		12	6.46	0.20	0.33	0.82	0.15	0.19	3.35	0.41	0.43	1.98	0.03	0.08	0.31	−0.39	−0.37
		17	7.03	0.25	0.37	0.97	0.20	0.24	3.72	0.56	0.57	2.15	0.05	0.10	0.18	−0.56	−0.54
Practical allocations																	
Vouchers		7	3.81	1.14	1.24	0.60	0.33	0.35	1.78	0.50	0.51	1.17	0.40	0.44	0.27	−0.09	−0.07
		12	6.25	1.25	1.34	1.11	0.43	0.46	2.90	0.63	0.64	1.92	0.44	0.48	0.32	−0.25	−0.24
		17	7.12	1.30	1.39	1.38	0.53	0.55	3.33	0.74	0.75	2.18	0.45	0.49	0.22	−0.42	−0.41
Vouchers, rural		7	3.61	1.12	1.21	0.72	0.39	0.41	1.51	0.43	0.44	1.13	0.39	0.44	0.25	−0.09	−0.07
		12	5.82	1.18	1.27	1.29	0.52	0.54	2.43	0.50	0.51	1.81	0.41	0.45	0.29	−0.26	−0.24
		17	6.67	1.22	1.31	1.63	0.64	0.67	2.78	0.57	0.59	2.07	0.43	0.47	0.19	−0.43	−0.41

Each cell reports the effect of subsidy policies, in the lower, focal, and upper equilibrium, for a given subsidy amount. Impacts represent the difference in these bounds. Implemented subsidy represents the negative of the impact of removing the subsidy from the baseline. For other policies the impact is of adding the counterfactual subsidy from the equilibrium where the implemented subsidy is removed. Utility and revenue reported in 2005 U.S. Dollars, discounted at a monthly rate of $\delta \approx 0.9945$. Consumer surplus includes the surplus utility each individual receives from the call model through May 2009, minus the cost of holding a handset from the time of adoption until May 2009. Rural accounts are defined as those whose modal tower is more than 15 km from a city.

implemented subsidy results from wider bounds in subsidy recipients' idiosyncratic value for phones.[24]

This is a setting with multiple equilibria. If the subsidy program also made individuals more optimistic about others' adoption (even those not directly or indirectly affected by the subsidy), then it could shift adoption into a more optimistic equilibrium. In that case, the impacts we report here would be an underestimate of the policy's effects.[25]

8. Conclusion

This project studies a mobile phone handset subsidy in Rwanda. Using transaction records after the subsidized handsets were activated, we trace how subsidized handsets are allocated and used across space. We use a structural model to simulate the effect of the enacted subsidy program, and compare it to counterfactual dispersals.

This project demonstrates how passively collected data can be used to evaluate the impact of programs and simulate the effect of counterfactual policies. This approach can also be combined with survey or experimental methods. For example, incentivized willingness to pay experiments could reduce uncertainty about subsidy recipients' idiosyncratic value of having a phone. Carefully designed experiments may also be able to shed light on how individuals form beliefs, and thus provide some guidance on which equilibrium a society is likely to coordinate on. However, it may be difficult to replicate the factors that influence coordination at a society scale using small experiments.

The effectiveness of different subsidy schemes will depend on the structure of spillovers. Spillovers in mobile phone adoption are mostly first order (contacts) and second order (contacts' contacts), and are bounded (there is only so much one person values speaking with each particular person). In contrast, for some goods, one node may induce substantial higher order spillovers (one idea, piece of viral content, or superspreader event could substantially affect the entire network). For these, global network structure is likely to be more important in targeting. The effect of a subsidy also depends on whether people are close to the margin of adopting (in which case small, dispersed interventions can tip the network), or there are clusters of people far from the margin (in which case larger, clustered interventions are likely to be more successful). A combination of empirical models of networks and reasoning about the structure of spillovers can inform the design most appropriate for a particular type of good.

This digital approach to impact evaluation can be applied in the increasing array of goods and services that are mediated by digital networks—such as digital credit and pay-as-you-go solar. This data and analysis can make policies and behaviors legible to centralized authorities (Scott, 1998), which can improve policymaking, but entails new risks of surveillance. Managing these risks will require society to have deep, informed conversations about what can be measured, and how it should be used.

[24] We allow targeting the subsidy only to eventual adopters. Note that the government could also subsidize individuals who do not appear in the data, but we are unable to empirically evaluate the impact of doing so. This could bias our results downward (if nonadopters are good candidates for subsidization) or upwards (if the included adopters are better candidates for subsidization, but a government would not have been able to identify them).

[25] One could bound the combined effect by reporting the difference between a pessimistic equilibrium without the policy and an optimistic equilibrium with it, for given values of $\boldsymbol{\eta}$; thanks to an anonymous referee for making this point.

Disclosures and data availability

This work was supported by the Stanford Institute for Economic Policy Research, United States of America through the Shultz Fellowship in Economic Policy and the JP Morgan Chase Research Assistant Program at Brown University. During part of the writing of this paper, Björkegren was supported as a Visiting Researcher at Microsoft Research.

Data was provided by the anonymous telecom partner in Rwanda as well as publicly available sources. Data from the first source was provided under the condition that it be kept confidential, so we are unable to provide that data.

Appendix

Adopting a phone may transform an individual's social network—they may keep in touch with friends or family living further away, for example. We uncover the communication graph after any transformation associated with adoption: the graph conditional on phone ownership. The inference in this paper remains valid as long as any such transformation is anticipated and coincides with adoption.

References

Alatas, V., Banerjee, A., Hanna, R., Olken, B.A., Purnamasari, R., Wai-Poi, M., 2016. Self-targeting: Evidence from a field experiment in Indonesia. J. Polit. Econ. 124 (2), 371–427.

Alatas, V., Banerjee, A., Hanna, R., Olken, B.A., Tobias, J., 2012. Targeting the poor: Evidence from a field experiment in Indonesia. Amer. Econ. Rev. 102 (4), 1206–1240.

Athey, S., Eckles, D., Imbens, G.W., 2018. Exact p-values for network interference. J. Amer. Statist. Assoc. 113 (521), 230–240, Publisher: Taylor & Francis _eprint: https://doi.org/10.1080/01621459.2016.1241178.

Banerjee, A., Chandrasekhar, A.G., Duflo, E., Jackson, M.O., 2013. The diffusion of microfinance. Science 341 (6144), 1236498.

Banerjee, A., Chandrasekhar, A.G., Duflo, E., Jackson, M.O., 2019. Using gossips to spread information: Theory and evidence from two randomized controlled trials. Rev. Econom. Stud. 86 (6), 2453–2490.

Beaman, L., BenYishay, A., Magruder, J., Mobarak, A.M., 2021. Can network theory-based targeting increase technology adoption? Amer. Econ. Rev. 111 (6), 1918–1943.

Björkegren, D., 2019. The adoption of network goods: Evidence from the spread of mobile phones in Rwanda. Rev. Econom. Stud. 86 (3), 1033–1060.

Breza, E., Chandrasekhar, A.G., McCormick, T.H., Pan, M., 2020. Using aggregated relational data to feasibly identify network structure without network data. Amer. Econ. Rev. 110 (8), 2454–2484.

Burton, M., Macher, J., Mayo, J.W., 2007. Understanding participation in social programs: Why don't households pick up the lifeline? B.E. J. Econ. Anal. Policy 7 (1).

Chin, A., Eckles, D., Ugander, J., 2021. Evaluating stochastic seeding strategies in networks. Manage. Sci. Publisher: INFORMS.

Cohen, J., Dupas, P., 2010. Free distribution or cost-sharing? Evidence from a randomized malaria prevention experiment. Q. J. Econ. 125 (1), 1–45, Publisher: Oxford University Press.

Cunha, J.M., De Giorgi, G., Jayachandran, S., 2019. The price effects of cash versus in-kind transfers. Rev. Econom. Stud. 86 (1), 240–281.

Domingos, P., Richardson, M., 2001. Mining the network value of customers. In: Proceedings of the Seventh ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. In: KDD '01, Association for Computing Machinery, San Francisco, California, pp. 57–66.

Farr, T.G., Rosen, P.A., Caro, E., Crippen, R., Duren, R., Hensley, S., Kobrick, M., Paller, M., Rodriguez, E., Roth, L., Seal, D., Shaffer, S., Shimada, J., Umland, J., Werner, M., Oskin, M., Burbank, D., Alsdorf, D., 2007. The shuttle radar topography mission. Rev. Geophys. 45 (2).

Feld, S.L., 1991. Why your friends have more friends than you do. Am. J. Sociol. 96 (6), 1464–1477, Publisher: University of Chicago Press.

Garbacz, C., Thompson, H.G., 1997. Assessing the impact of FCC lifeline and link-up programs on telephone penetration. J. Regulatory Econ. 11 (1), 67–78.

Garbacz, C., Thompson Jr., H.G., 2005. Universal telecommunication service: A world perspective. Inf. Econ. Policy 17 (4), 495–512.

Ghose, D., 2018. Chhattisgarh: 50 lakh people to get smartphones under sanchar kranti yojana. Indian Express.

GSMA, 2013. Universal Service Fund Study. Technical report, GSM Association.

Higgins, S., 2019. Financial technology adoption. Working Paper.

Hjort, J., Poulsen, J., 2019. The arrival of fast internet and employment in Africa. Amer. Econ. Rev. 109 (3), 1032–1079.

Jack, W., Suri, T., 2014. Risk sharing and transactions costs: Evidence from Kenya's mobile money revolution. Amer. Econ. Rev. 104 (1), 183–223.

Jarvis, A., Reuter, H., Nelson, A., Guevara, E., 2008. Hole-filled seamless SRTM data V4. CIAT.

Jensen, R., 2007. The digital provide: Information (technology), market performance, and welfare in the south Indian Fisheries sector. Q. J. Econ. 122 (3), 879–924.

Kempe, D., Kleinberg, J., Tardos, E., 2003. Maximizing the spread of influence through a social network. In: Proceedings of the Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. In: KDD '03, ACM, New York, NY, USA, pp. 137–146.

Kim, D.A., Hwong, A.R., Stafford, D., Hughes, D.A., O'Malley, A.J., Fowler, J.H., Christakis, N.A., 2015. Social network targeting to maximise population behaviour change: A cluster randomised controlled trial. Lancet 386 (9989), 145–153.

New Times, 2008. Residents get heifers, telephone handsets.

Perkins, J.M., Subramanian, S.V., Christakis, N.A., 2015. Social networks and health: A systematic review of sociocentric network studies in low- and middle-income countries. Soc. Sci. Med. (1982) 125, 60–78.

Scott, J.C., 1998. Seeing Like a State: how Certain Schemes to Improve the Human Condition Have Failed. Yale University Press.