CENTRE FOR THE STUDY OF AFRICAN ECONOMIES

CSAE Working Paper WPS/2018-18

Access & Excess - The Effect of Internet Access on the Consumption Decisions of the Poor

Pieter Joseph Sayer*

2nd November 2018

Abstract

This paper tests for the existence of a causal relationship between internet access and the consumption decisions of the poor. In particular, we develop theoretical models that provide two testable hypotheses. Firstly, that, by better aligning beliefs with reality regarding the value of future income, information attained via the internet will either increase or decrease one's propensity to consume, depending on the nature of their occupation. Second, that, by heightening one's concern for their relative status, access to social media via the internet will cause users to increase their conspicuous consumption. The analysis exploits a natural experiment, the 2015 roll out of a free internet provision program (Free Basics) in Colombia, as a source of exogenous variation in potential internet access. By taking advantage of a unique geo-coded dataset of telecommunications pylons in Colombia this paper is able to compare households covered by the free internet provision program to those who are not. We estimate a number of 2SLS instrumental variable and difference-in-differences specifications on rich data from the Encuesta Longitudinal Colombiana (ELCA) panel survey. Our estimates provide tentative evidence in support of our second hypothesis as we estimate a local average treatment effect (LATE) of a 12.2 percentage point increase in the share of conspicuous consumption in the overall consumption bundle as a result of internet access. We find no evidence of a statistically significant LATE on overall consumption. Our difference-in-differences, intention-to-treat (ITT), estimates for both outcomes are small and insignificant, however this could be a result of a lack of power given the low take-up of the internet provision program.

*Email: pietersayer@hotmail.com. I am grateful for excellent supervision from Rossa O'Keefe-O'Donovan and for superb comments and support from Martin Montenegro & Jorge Cuartas.

UNIVERSITY OF OXFORD

Centre for the Study of African Economies

Department of Economics · University of Oxford · Manor Road Building · Oxford OX1 3UQ

T: +44 (0)1865 271084 · F: +44 (0)1865 281447 · E: csae.enquiries@economics.ox.ac.uk · W: www.csae.ox.ac.uk

1 Introduction

The discourse among the development community largely takes for granted that the diffusion of ICT is a transformative and ultimately desirable objective; yet evidence in favour of such a supposition is sparse. The literature shows any evidence of a macro relationship between connectivity and growth to be inconclusive and causal work of any kind is a rarity (Friederici, Ojanpera & Graham, 2017; World Bank, 2016). While none of this is to say that internet access definitely does not or cannot have the kinds of positive effects that one might intuit, there is an apparent mismatch between this lack of evidence and the seeming consensus in the policy sphere. There is a clear gap in the literature, particularly with regards to the mechanisms through which a relationship might take shape. Considering the billions of dollars being spent on expanding internet penetration across the developing world it seems pertinent to understand the impacts such projects might have and whether or not they represent the most efficient use of development funding.

This paper contributes to the understanding of the impacts of internet diffusion by testing for evidence of a causal relationship between internet access and the consumption decisions of the poor. We exploit a natural experiment, the roll out of Free Basics in Colombia, as a source of exogenous variation in potential internet access. Free Basics, a Facebook led initiative aimed at delivering free internet access to people in developing countries, is provided exclusively through the Tigo mobile network in Colombia. Accordingly, if one can assume that being covered by Tigo as opposed to other networks in Colombia is arbitrary, then the increase in access afforded by Free Basics' release on the 14th of January 2015 can be seen as exogenous.

There are a number of potential mechanisms through which a shock in access to the internet might change consumption patterns. This paper will focus on two hypotheses which can be tested in our data. The first is that access to information may correct beliefs about future payoffs. Such information could come through use of government websites which provide details of future social transfers and on the returns to human and working capital or through the wisdom of crowds available through social media. The second, more novel mechanism, runs through social media; people's consumption may be shifted towards more conspicuous consumption as they gain increased utility from certain luxury goods and activities that are deemed 'share-worthy'. One could hypothesise this to come through the social validation and status available through 'showing off' on social media - think "likes" on Facebook. The theory behind these mechanisms will be expanded upon in Section 3.

We estimate, among other specifications, a 2SLS instrumental variable estimate of the LATE as well as a difference-in-differences estimate of the ITT using Colombian panel data. We find tentative evidence of a causal relationship between internet access and conspicuous consumption in our statistically significant estimate for the LATE of a 12.2 percentage point increase. We consider this evidence as only tentative because it is not confirmed by our tests of robustness and because we cannot completely rule out a violation of our exclusion restriction. We do not find convincing evidence in support of our belief correction hypothesis.

While the analysis marks the first steps towards understanding the causal impact of internet diffusion on consumption behaviour, there does exist a growing literature considering the role of the internet, and ICT more broadly, on a spectrum of outcomes. The greater literature on ICT and internet diffusion largely revolves around two themes: cross-country analysis of their relationship with economic growth and studies considering their effect on enhancing human capital accumulation. Neither of these literatures is by any means in a state of consensus despite the strong priors of a positive relationship that most people would have. This disconnect again points to the need for more rigorous work on all aspects of internet diffusion.

The body of work concerning human capital accumulation shows the identification of a causal relationship between computer assisted learning and test scores to be hugely context dependent. RCTs have consistently identified a causal positive relationship in the contexts of China and India however they have consistently failed to do so in the Latin American context (Banerjee et al., 2007; Barrera-Osorio & Linden, 2009; Beuermann et al., 2015; Glewwe & Muralidharan, 2016). Evidence regarding the role of internet in this mechanism is also found wanting and the only study to consider internet access, that of Cristia, Czerwonko and Garofalo (2014), fails to find an effect using a difference-in-differences approach in Peru. With regards to growth, there does seem to be a slight majority of the high-quality papers showing a positive causal effect in developed countries using instrumental variable approaches; however, the literature on developing countries is less convincing (Czernich et al., 2011; Galperin & Viecens, 2017).

Perhaps the most relevant work with regards to this paper's topic is Jensen (2007) which identifies a positive causal link between mobile phone usage and societal welfare in a developing context. Jensen finds that the effect runs through a reduction in informational asymmetries. If one sees mobile internet as a means of enhancing this effect through even greater scope for information sharing, then one could see Jensen (2007)'s findings as supportive of this paper's hypothesis; as people become better informed they will make better consumption decisions. This hypothesis also has theoretical groundings in Kremer, Mansour and Perry (2014) which provides a formal model for the 'feedback effect' though which agents may harness the wisdom of crowds via the internet and social media in particular. A third source of direct support from the literature is that of Lohmann (2015); the paper uses European panel data to identify a positive link between internet access and individual positionality. It is through this mechanism, of an increase in people's material aspirations and in their concern for their relative status, through which this paper hypothesises that access to social media will impact consumption decisions.

Upon review of the literature it has become evident that not only is there a clear gap in the literature with regards to the subject of this paper, but that there is very little rigorous work concerning the effects of internet access at all. The novelty of our topic, approach and data allow us to make a significant contribution towards closing this gap.

The remainder of this paper is structured as follows. In Section 2 we outline the Free Basics project and highlight the cost barriers it overcomes in the Colombian context. Section 3 presents the theoretical framework upon which the empirical analysis is based. In Section 4 we set out the empirical strategy for testing the specified relationships and outline any assumptions made as well as any evidence regarding their validity. Section 5 reports the associated findings and their robustness. Finally, Section 6 discusses our results, their policy implications and avenues for further research.

2 Context

2.1 Free Basics

Facebook's Free Basics program is aimed at removing the cost barrier that prevents millions of people in developing countries from accessing the internet. Facebook have identified that while over 85% of the world's population live in areas with mobile network coverage, only around 50% of people actually use the internet (Facebook, 2018; World Bank, 2017). This huge disparity implies that, for the remaining 35%, most of whom live in the developing world, there must exist barriers beyond coverage and cost is likely one of them. Through its partnerships with network service providers in developing countries, Free Basics effectively fully subsidises mobile data used on their platform. They claim that via this mechanism they have been directly responsible for bringing over 25 million people online who otherwise would not be (Facebook, 2018). The platform gives users access to a selection of websites that varies by context; in Colombia's case the websites available include Facebook, Wikipedia and a number of informative websites run by the government and NGOs. The websites provide users with the means to learn about government programs as well as a plethora of information that should theoretically help users make better decisions in a number of arenas, including consumption spending. There are a number of websites explicitly focused on sharing information on all aspects of agricultural activities which should be useful in informing the investment and saving decisions of the rural population in particular. A list and description of selected websites is available in Table 6 of the Appendix.

2.2 Mobile Internet in Colombia

Free Basics was released in Colombia on the 14th of January 2015 and until 2018, which is outside the range of our data, it was exclusively available through the Tigo mobile network. In 2016, Tigo captured 18% of the country's competitive mobile network market; its main rivals, Claro and Movistar captured 55% and 25% respectively (GSMA, 2016). While these three rivals combine to provide widespread 3G coverage to over 92% of the population, only 52% of Colombians actually boast a mobile internet connection; a prime example of the disparity between coverage and adoption that Free Basics seeks to resolve (GSMA, 2017). Cost is frequently cited as the next major roadblock in extending access and in a 2016 survey 52% of non-users named affordability as a barrier to their connection (GSMA, 2017). Mobile internet is the medium of choice in Colombia with 79% of internet users reporting to have done so through a mobile device; this is likely a result of the relatively low cost of mobile handset in comparison to other devices. Despite it being the cheapest means of access, it is still by no means cheap in this context; the most basic smartphones can be bought for 50,000 Colombian Pesos (39$ PPP) but data costs can be prohibitive[1]. The median cost of 100MB per week would account for over 10% of the monthly income of a person living below the Colombian poverty line of 241,673 Pesos (187$ PPP) per person per month (DANE, 2017)[2]. With 28% of Colombians defined as poor by this measure it seems clear that data costs are a constraint for a large subsection of the population (DANE, 2017).

3 Theoretical Framework

3.1 Adjusting Beliefs

The first hypothesised mechanism we consider formally is that of information provided via the internet altering people's beliefs about the future. Consider a utility function defined over a two-period life cycle where the agent can both save for the future and borrow from the future in period 1 but can do neither in period 2. The

[1]Table 7 of the Appendix provides a summary of handset and data costs by provider.

[2]100MB is equivalent to around 4 hours of browsing time which this paper defines as the minimum time requirement for an individual to make productive use of the internet (Lotta, 2017).

model is set up as a two-period life cycle for the simplicity of notation. It can be extended to an $n > 2$-period model without loss of generality but this would not significantly alter or add to the interpretation. The agent is assumed to invest any savings made and has naive beliefs with respect to both their future income and the rate of return on their period 1 investments. The agent maximises the following utility function:

$$U(c_1, c_2) = lnc_1 + \beta lnc_2 \tag{1}$$

subject to the lifetime budget constraint $c_1 + \frac{c_2}{(1+r)} = y_1 + \frac{y_2}{(1+r)}$ where $y_2 = y_1\delta$ and r represents the rate of return on investment. One could see r as capturing gains available from investment opportunities and y_2 as capturing gains available from all other sources of income that are independent of savings such as job-market activities and social transfers. Accordingly, r is most salient for entrepreneurs or land-owning farmers whereas y_2 would be more salient for employees. β is a discount factor that captures the agent's time preferences s.t. $\beta \in [0, 1]$ and δ is defined as $\delta \in \mathbb{R}^+$ where the agent has some belief about δ, δ^B. Similarly, r is defined as $r \in \mathbb{R}^+$ where the agent has some belief about r, r^B. With the true parameters, δ and r, unknown in period 1, the agent will maximise utility based on their naive beliefs about their values, acting is if they were the true parameters. This maximisation gives the following first order conditions:

$$FOC_{c_1} = \frac{1}{c_1} - \lambda$$
$$FOC_{c_2} = \frac{\beta}{c_2} - \frac{\lambda}{(1 + r^B)}$$

Solving for consumption, we get $c_1^* = \frac{y_1}{(1+\beta)}\left(1 + \frac{\delta^B}{(1+r^B)}\right)$. In period 2, y_2 and r are revealed such that the agent's period 2 consumption is determined by the true δ and r. c_2 is therefore equal to the income in the second period plus any savings from the first period and minus any income borrowed in the first period.

$$c_2^* = (y_1 - c_1^*)(1 + r) + y_1\delta \tag{2}$$

$$= \frac{y_1}{1+\beta}\left((1 + r + \delta)\beta + \delta - \delta^B\left(\frac{1+r}{1+r^B}\right)\right) \tag{3}$$

By substituting the marshallian demands into the utility function we get the utility function in terms of δ, δ^B, r & r^B as shown in equation (4).

$$\begin{aligned} U(\delta^B, \delta, r^B, r) = ln\left(\frac{y_1}{1+\beta}\left(1+\frac{\delta^B}{1+r^B}\right)\right) \\ + \beta ln\left(\frac{y_1}{1+\beta}\left((1+r+\delta)\beta+\delta-\delta^B\left(\frac{1+r}{1+r^B}\right)\right)\right) \end{aligned} \tag{4}$$

After differentiating equation (4) w.r.t. δ^B to give equation (5) one can rearrange to show that $\frac{dU}{d\delta^B} = 0$ only when $\delta^B = \delta$. Equation (6) then shows us that $\frac{d^2U}{d\delta^B d\delta^B} < 0$ as each individual component is defined as greater than or equal to 0. The combination of these two results shows that utility is maximised w.r.t. δ^B at $\delta^B = \delta$ or in other words, that utility is a decreasing function of $|\delta - \delta^B|$. Similarly, by differentiating w.r.t. r^B in place of δ^B one can show that utility is also a decreasing function of $|r - r^B|$. These results highlight the costliness of mistaken beliefs.

$$\frac{dU}{d\delta^B} = \frac{1}{(1+r^B+\delta^B)} - \frac{\beta(1+r)}{(1+r^B)((1+r+\delta)\beta+\delta) - \delta^B(1+r)} \tag{5}$$

$$\frac{d^2U}{d\delta^B d\delta^B} = -\frac{1}{(1+r^B+\delta^B)^2} - \frac{\beta(1+r)^2}{((1+r^B)((1+r+\delta)\beta+\delta) - \delta^B(1+r))^2} \tag{6}$$

If one considers the gaps between the agent's beliefs and the true parameters to be functions of information regarding δ and r, i^δ and i^r respectively, such that $\frac{d|\delta-\delta^B|}{di^\delta} < 0$ and $\frac{d|r-r^B|}{di^r} < 0$ then an increase in information via the internet should lead to a shift towards a consumption allocation associated with a higher level of utility. Information regarding future income may come through information about subsidies or social benefits found on the government websites or through the wisdom of crowds available through social media for example. With regards to return on investment, Free Basics provides a myriad of sources of information on education, weather, health and agriculture which should reduce uncertainty regarding investment in both human capital and material inputs. The anecdotal evidence of the outcomes of those who have engaged in such investments in the past available on social media should also play a clarifying role. It is most likely that without said information, agents underestimate the opportunity cost of not investing and as well as the value of alternative income sources and so $r^B < r$ and $\delta^B < \delta$. As $\frac{dc_1^*}{d\delta^B} > 0$ and $\frac{dc_1^*}{dr^B} < 0$ the resultant correction should have opposing effects on consumption and so the direction of any change we see in the data would show which effect is dominant and in turn where the information asymmetry between the connected and the unconnected is greatest.

In terms of our empirical application of the model it is important to address any concerns over which period the agent is in when we actually observe them in the data. One such concern would be any income effects that may arise from the fact that both the true r and δ, and any beliefs concerning them, may be increasing functions of the aforementioned information; this would lead to increases in both current and future consumption. It may also be the case that those who initially save more as a result of internet access end up consuming more in gross terms as they reap the benefits of their increased investments. This kind of behaviour would lead to us misinterpreting results after a certain post-treatment period. To account for this issue we condition on income in the empirical work such that we consider given one's income, how much do they consume. By extending the model to an $n > 2$-period life cycle, one can show that, after conditioning on income, the directions of the predictions concerning period 1 consumption are appropriate for any period in which there is a subsequent period over which the agent holds beliefs. Period n, can be seen as representing the agent's final earning period and so, given the infeasibility of retirement in our context, and the unpredictability of death, the period 1 predictions will, for the majority of the population, pertain to any observable period. That said, one might predict that the magnitude of the effect may be a decreasing function of the probability that the current period is in fact the final period. In this way the period in which we observe the agent should only be a concern in terms of the magnitude of the effect, not the direction.

3.2 Showing Off

The second mechanism we look to test in the data is that of an increased propensity to 'show off' as a result of the increased audience available via social media. One could see this mechanism working to some degree through the relationship between one's concern for their relative status and their subjective well-being. Social media is notorious for the positive selection bias users have with regards to the nature of their posts and its normalising of images of a materially successful lifestyle. As a result, viewers of social media naturally face downward pressure on their perception of their social status relative to that of the peers (Lohmann, 2015). The validation and status available through the posting of 'share-worthy' goods and services is a means through which one can counteract this downward pressure. Consider a utility function defined over the consumption of share-worthy (SW) and not-share-worthy (NSW) goods and services.

$$U(SW, NSW) = (1 + f(\theta))lnSW + ln(NSW - \tau) \tag{7}$$

subject to the budget constraint $m = P_{SW}SW + P_{NSW}NSW$. Where θ represents the potential social exposure of a good or service and is defined as greater than or equal to 0 and as enhancing the utility derived from SW such that $\frac{dU}{df(\theta)} > 0$. We assume $f(\theta)$ is increasing, concave and continuously differentiable. τ is some strictly positive subsistence parameter that must be met such that NSW$> \tau$. To solve the agent's utility maximisation problem we derive the following first order conditions:

$$FOC_{SW} = \frac{1 + f(\theta)}{SW} - \lambda P_{SW}$$
$$FOC_{NSW} = \frac{1}{(NSW - \tau)} - \lambda P_{NSW}$$

By solving for SW and taking the first derivative one can determine if the optimal share of SW goods and services is an increasing or decreasing function of $f(\theta)$:

$$\frac{dSW^*}{df(\theta)} = \frac{m - \tau P_{NSW}}{P_{SW}(2 + f(\theta))^2} \tag{8}$$

As all the components of $\frac{dSW^*}{df(\theta)}$ are individually defined as greater than or equal to 0 then it is clear that $\frac{dSW^*}{df(\theta)} > 0 \;\forall\; m > \tau P_{NSW}$; once the value of the subsistence requirement is covered, the share of share-worthy goods in the optimal consumption bundle increases. By solving for SW instead of NSW in equation (8) one can also show that $\frac{dNSW^*}{df(\theta)} < 0 \;\forall\; m > \tau P_{NSW}$ by following the same process. In line with this mechanism we test whether or not the share of 'share-worthy' goods and services in the agent's consumption basket increases as a result of access to social media; that is to say, is $f(\theta) > 0$.

4 Empirical Methodology

If Free Basics significantly increases internet access our theoretical models suggest two testable hypotheses:

1. The total consumption of those covered by Tigo will be different from those who are not, conditional on income[3].

2. The share of 'share-worthy' goods and services in the consumption bundle of those covered by Tigo will be higher than that of those not covered by Tigo[4].

4.1 Data

The data used to test these hypotheses comes from two sources, the Encuesta Longitudinal Colombiana (ELCA) and the Ministry of Technology, Information and Communication in Colombia (MinTIC). The ELCA is a geo-coded longitudinal household survey study that consists of a baseline round in 2010 and follow ups in both 2013 and 2016. The survey provides detailed consumption information on a sample that contains over 10,000 households and is nationally representative of those in social strata 1 to 4[5]. The ELCA contains information on over 72 categories of consumption and so the analysis will consider the aggregate value spent on these categories as well as the share of total consumption spent upon the categories deemed to be 'share-worthy'. The 72 types of consumption are reported for varying time intervals and so we standardise values to represent the expenditure of an average month; yearly values are divided by twelve and so on. Our definition of 'share-worthy' expenditure includes categories such as personal grooming, clothing, entertainment and vacations[6]. In line with our focus on those for whom cost is a barrier to internet access, we only consider the 7,000 households in the ELCA that reside in social strata 1 and 2; those defined as 'lower-class' by the Colombian government (DANE, 2018)[7].

[3] Unless relative reductions in the information deficits with regards to return on investments and with regards to saving-independent future income sources exactly offset each other.

[4] While we hypothesise a positive value we still run two-sided t-tests as to not impose our priors on the test.

[5] The Colombian government classifies neighbourhoods on a scale of 1 to 6 based on a number of socioeconomic indicators; 1 represents the poorest neighbourhoods and 6 the richest.

[6] A full list of the goods and services deemed 'share-worthy' is available upon request.

[7] We also exclude the top 5% of earners from the sample as they earn over the equivalent of 2,000 2016 PPP dollars per month which we again, don't consider to be poor.

The collection of the 2016 round of the ELCA took place in January of that year and so, with Free Basics release date being in January of 2015, our analysis pertains to the effects of Free Basics over a one-year horizon.

Figure 1: Maps of Sample Municipalities and Telecommunications Towers

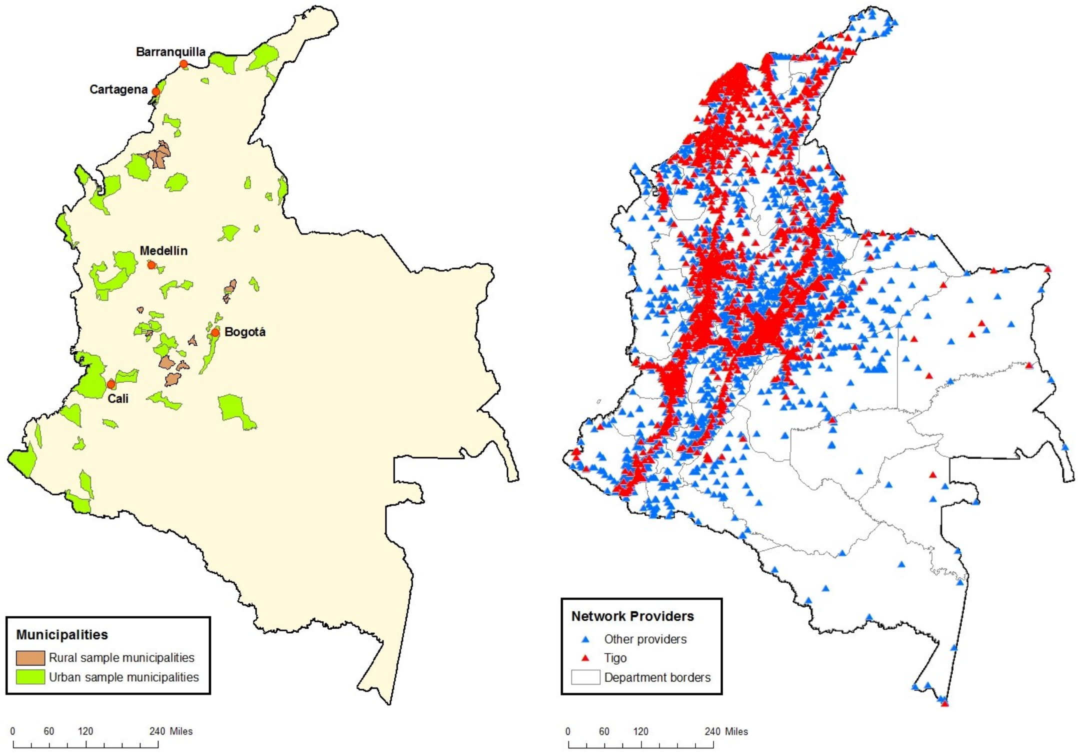

The unique pylon data provided by MinTIC includes information on the coordinates, the operating provider and the transmission technology used for every communications pylon in Colombia. These details allow us to uniquely identify Tigo towers and to accurately estimate the towers' coverage range. This data set has not, to our knowledge, been used for any rigorous analysis to date. The GPS coordinates for the ELCA households will be combined with GPS coordinates for the communications pylons to construct distance variables which will be used as indicators of internet coverage and potential access.

4.2 Empirical Approach

In the analysis we look to estimate both the intent-to-treat (ITT) and the local average treatment effect (LATE) to provide an indication of the program's effects at both the aggregate and individual level. With that said, we treat the LATE effect as our main result for two reasons. Firstly, this paper is focused on the effects of having internet access on an individual's consumption decisions and so the LATE is a more appropriate metric in this way. Second, we have to keep in mind that, with only 24% of our 'treated' subsample reporting access to the internet after treatment, take up of the Free Basics initiative was clearly not huge and so we may be underpowered to identify an estimate for the ITT. As a result we provide the ITT as a secondary result for the benefit of policy makers who are likely more concerned with the net effects of a program like Free Basics.

Whether or not a household is covered by Tigo is, on its own, likely endogenous due to a number of selection criteria issues that surely come into play when networks plan where to cover. In order to combat this, we condition on the fact that the household has network coverage at all; making the assumption that, as a result of the competitive nature of Colombia's telecommunications market, the networks have a similar selection criterion for providing coverage. The analysis will therefore use the distance of households from Tigo telecommunication pylons as an instrument for the 'treatment', the availability of cheaper internet, conditional on being covered by at least one provider.

One of the key decisions in the analysis pertains to the definition of coverage; where do we draw the line for the radius of coverage provided by a telecommunications tower. This paper makes such a decision on the basis of three sources; expert opinion from an official at MinTIC, a report from a global provider of wireless communications infrastructure and the data. A tower's coverage range depends on a number of factors such as capacity, topology and technology and so there is no one-size-fits-all range. As such, this paper looks to provide a best estimate with the maximum number of degrees of freedom possible within the given time constraints. The main differentiation we allow for is between rural and urban areas as we have a clear distinction between these zones in the ELCA data. American Tower, a global provider of wireless communications infrastructure, and MinTIC are in general consensus that ranges can vary from as little as 500m in very densely populated urban areas to over 5km in sparse rural regions (American Tower, 2014). Our data is relatively consistent with these estimates as we see quite clear cut-offs in reported internet access for those further than 4.8km from a tower in the rural zone and for those over 3km from a tower in the urban zone[8]. These cut-offs are very close to the proposed ranges of 2.5km in suburban areas, an appropriate definition of the peripheral urban zones inhabited by Colombia's poor, and 5km for rural areas and so we proceed with the ranges defined in the data (American Tower, 2014).

While the ranges of 3G and 4G towers are not hugely different, those carrying 2G technology do tend to have a significantly larger range which could suggest heterogeneity in the accuracy of our defined cut-off. We find however that our results are robust to the removal of 2G towers which suggests technology is seemingly not the limiting factor with regards to coverage and so our lack of accounting for it shouldn't bias our results. Further research might consider adding more degrees of freedom with regards to technology and topology to give more accurate cut-off estimates.

4.2.1 Instrumental Variable

Our main specification considers a 2SLS instrumental variable (IV) approach which exploits a question in the ELCA concerning whether or not the household possesses some form of internet connection. We instrument for this reporting of a connection with coverage by Tigo in the post-treatment period, 2016, as well as its interaction with distance from a Tigo tower. We add the interaction with distance from a tower as we hypothesise that quality of coverage is a decreasing function of distance to a tower and so those closer to a tower should be more likely to take up treatment. The first stage is displayed in equation (11) and the structural equation is depicted in equation (12).

$$internet_{ht} = \alpha_0 + \alpha_1 Tigo_h + \alpha_2 T_{ht} + \alpha_3 (T \text{ x } d)_{ht} + \mathbf{X}_{ht} + \gamma_h + \rho_t + \epsilon_{ht} \quad (9)$$

$$y_{ht} = \beta_0 + \beta_1 Tigo_h + \beta_2 \widehat{internet}_{ht} + \mathbf{X}_{ht} + \gamma_h + \rho_t + \varepsilon_{ht} \quad (10)$$

where h indexes household and t indexes time. y_{ht} is the consumption outcome considered which is either the log of the total value of consumption (log(Consumption)) or the share of the total value of consumption spent on 'share-worthy' goods (SW

[8] Graphs of internet access against distance from a tower are available upon request.

share). γ_h and ρ_t represent household and time fixed effects respectively. $\mathbf{X}_{ht}$ is a vector of covariates including coverage, household characteristics and characteristics of the head of the household; a full list is available in Table 4 of Section 5. $Tigo_h$ is a dummy variable with a value of 1 if the household is within the defined radius of coverage of a Tigo tower and d is distance from a Tigo tower. T_{ht} is the treatment term which is equivalent to $Tigo$ x 2016_{ht}, an interaction term between $Tigo$ and a dummy for it being 2016, the post-treatment survey round. If our hypotheses are correct we would expect α_2 and α_3 to have countervailing signs in the first stage such that distance from a tower acts as a dampening effect on the treatment's ability to predict *internet*. We also expect the coefficients β_1 and α_1 to be 0 if our assumption of exogenous allocation of treatment is correct. $\widehat{internet}$ represents the partial of *internet* which covaries with our instruments, $Tigo$ x 2016 and $Tigo$ x 2016 x d. β_2 therefore gives the LATE of Free Basics on our two outcomes.

4.2.2 Difference-in-Differences

We also use a simple difference-in-differences (DID) strategy to estimate the ITT effect of Free Basics on our two outcomes.

$$y_{ht} = \beta_0 + \beta_1 Tigo_h + \beta_2 Tigo \text{ x } 2016_{ht} + \mathbf{X}_{ht} + \gamma_h + \rho_t + \varepsilon_{ht} \tag{11}$$

where y_{ht}, $\mathbf{X}_{ht}$, γ_{ht}, ρ_t, $Tigo_h$ and $Tigo$ x 2016_{ht} are as described in Section 4.2.1. Again we would expect the coefficient β_1 to be 0 if our assumption of exogenous allocation of treatment is correct. β_2 can be interpreted as the ITT effect of Free Basics on the outcome y_{ht}. The key assumption here is that, had Free Basics not been released, the outcomes' trends would have been the same for both those covered by Tigo and those covered by a different provider.

4.3 Robustness

Whether or not we can make the two key assumptions necessary in order to make meaningful causal inference based on this paper's results, those of validity and parallel trends, essentially boils down to whether or not we believe that treatment is essentially arbitrary. With the assumption that being covered by Tigo is exogenous conditional on coverage being so pivotal for the reliability of any inference based on any of this paper's results we consider six further variations on the identification strategy in case the assumption proves too stringent in the full sample. It may be the case that there are issues of market segmentation or legacies of natural monopolies that cause imbalance in unforeseeable ways and so we provide alternative strategies that should be more robust to such sources of bias. We consider five variations to provide robustness to our IV approach and a matching approach to provide robustness to our ITT estimate. We present the results of these methods in Section 5.2.

The first method we consider is to restrict the sample to those covered by one provider or less. The logic is that being covered by multiple providers can be seen as a signal of one's desirability as a client for the network providers which would make them a distinct group to those covered by one or no providers. One may see someone covered by a single provider as someone more likely to be covered not as a result of targeting but more so by chance due to their proximity to a targeted area. In this way one could see being covered by only Tigo as exogenous even without conditioning on coverage in a sample restricted to those covered by one provider or less if one can assume none of them are targeted by providers. That said, the assumption that those covered by one or no providers are identical in terms of targeting by providers may be too stringent and so the second variation exclusively considers households covered by one provider.

The three remaining methods are all variants of a fuzzy regression discontinuity design (RD). The RD allows us to loosen assumptions of balance across the whole sample and focus on balance of those within a set radius of the cut-off of coverage. The RD is fuzzy as there is no strict cut-off for treatment, we simply assume that the probability of being treated, $p(d_h)$, is increased when distance from a Tigo tower, d_h, is less than the defined cut-off. This method essentially relies on the assumption that those who live in such close proximity to each other are more likely to be comparable but other than that is no different from a standard 2SLS IV approach. As such, we simply consider the fuzzy RD designs as a means of comparing what should theoretically be more similar groups while remaining consistent with the IV methodology set out in Section 4.2.2. In this case we define the radius as 2km from the defined coverage range. We extend this methodology slightly in two ways. Firstly, due to the inexactness of a cut-off of coverage we include a 'donut' design in which those within a certain radius of the cut-off are excluded from the analysis (Barecca, Lindo & Waddell, 2015). The idea being that by extending the width of the cut-off we are more likely to correctly identify it. In this case we exclude households within 500m of the cut-off which leaves the sample including only those households where $500m < |D - d_h| \leq 2000m$[9]. In exchange for this improvement in identification there is of course a reduction in sample size and the need to extrapolate estimates across the 'donut' space. Secondly, we limit the RD to those who are covered by a provider other than Tigo such that we are only comparing those who can be seen to have been targeted by providers. This follows from the logic previously put forward concerning their being a distinct group.

Figure 2: Sampling Strategy Visualisations

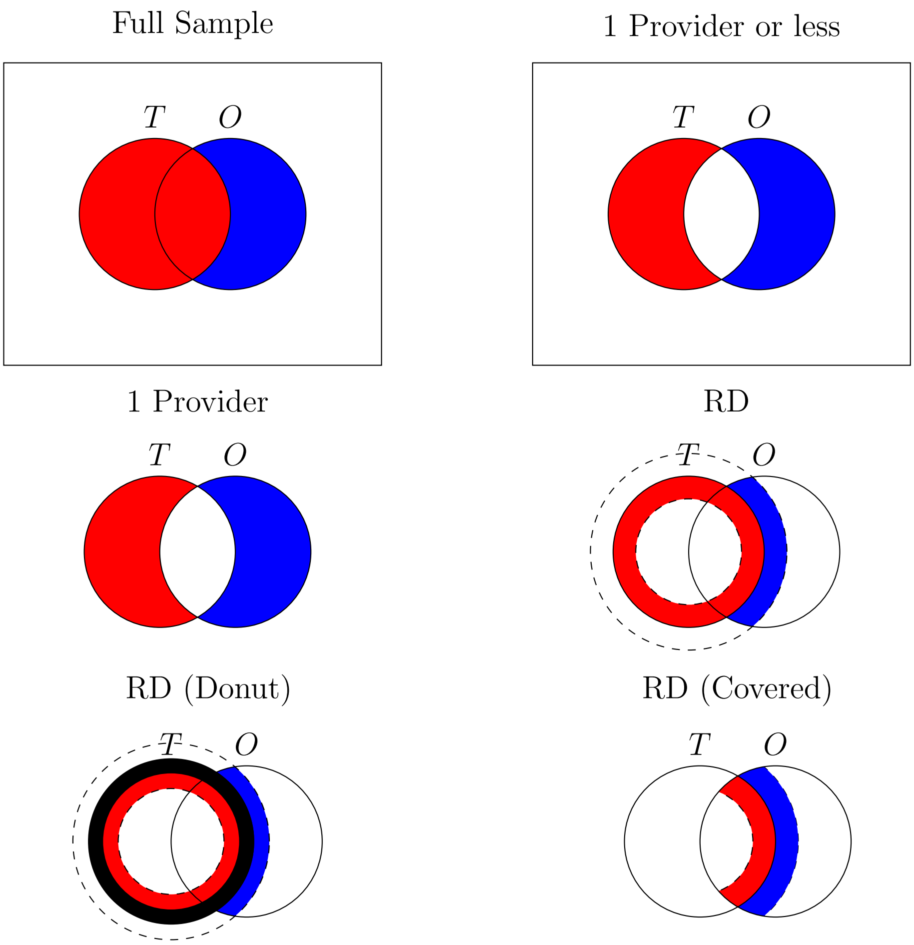

Figure 2 illustrates the sampling strategies more intuitively. By simplifying down to

[9]Our decisions to define the RD radius as 2km and the donut radius as 500m are admittedly rather arbitrary, our main consideration was maintaining sufficient power.

two types of coverage; that of Tigo (T) and that of one the four other providers (O), the interactions between the coverage ranges can be shown in the Venn diagrams below. The rectangles represents the full sample of households and within that there are households within the left-hand circle, T, who are covered by Tigo and households within the right-hand circle, O, who are covered by one of the other four providers. In the variations absent a rectangle the uncovered population is excluded or, in the case of the RD and RD (Donut) strategies, they are only included if within the radius from the cut-off defined by the RD design; this radius is represented by the dashed circles. Red areas are those covered by Tigo and included in the sample, blue areas are those covered by another provider (and not Tigo) that are included in the sample and the black area covers those excluded from the sample as part of the 'donut' strategy.

Table 1 brings context to the state of coverage in our sample as well as provides some indication as to the size of the blue and red areas in Figure 2. One main takeaway from Table 1 is that the vast majority of households exist in the area represented by the overlap between circles in the Venn diagrams. In the full sample 80% of covered households are covered by both Tigo and another operator and the average household is covered by more than two providers.

Table 1: Coverage Statistics

	Full Sample	<2 Providers	1 Provider	RD	RD (Donut)	RD (Covered)
Covered by Tigo	0.69	0.14	0.33	0.61	0.64	0.75
Covered by other	0.77	0.28	0.67	0.70	0.70	1.00
Covered	0.81	0.42	1.00	0.79	0.79	1.00
No of providers	2.29	0.42	1.00	1.78	1.85	2.42
Households	6546	1976	836	2672	2033	1870

Our final method is that of a matched difference-in-differences (MDID) approach. We have opted for Mahalanobis distance matching (MDM) using an Epanechnikov kernal to generate a control sample where households are weighted by the number of matches they make with households in the treatment group[10]. We use the generated treatment and control groups and the associated weightings to run the same DID specification outlined in Section 4.2.2. We choose MDM over the more frequently used propensity score matching (PSM) method to avoid the "propensity score paradox" outlined in King and Nielsen (2016)[11]. We use kernal matching as opposed to the nearest-neighbour method due to the greater flexibility it allows in weighting the control sample.

We would have also liked to perform a robustness check with regards to the mechanism of our first hypothesis, the relationship between internet access and total consumption. Our model suggested that the effect will flow in different directions for different sub-populations. We expect information regarding returns on investments to be more salient for farmers or entrepreneurs such that $\left|\frac{dc_1^*}{dr^B}\right| > \left|\frac{dc_1^*}{d\delta^B}\right|$ whereas we expect information regarding investment-independent income to be more salient for salaried employees such that $\left|\frac{dc_1^*}{dr^B}\right| < \left|\frac{dc_1^*}{d\delta^B}\right|$. This would entail that the self-employed would consume less as a result of treatment and salaried workers would consume more as a result of treatment. Unfortunately the data in the ELCA pertaining to the head of household's type of occupation is too incomplete and so this will have to remain as an avenue for future research.

4.4 Underlying Assumptions

With regards to the parallel trends assumption necessary for inference based on our DID results, we are limited in our ability to provide evidence either for or against it in that we only have three time periods to compare. The results of this comparison are included in Table 2 and we find that trends seem to have been somewhat parallel between 2010 and 2013 for the full sample. What is striking is that consumption seems to be extremely balanced in this regard which is promising for our results concerning our first hypothesis. The 'share-worthy' share's trend is less encouraging as is that of income, despite its statistical insignificance. That said, when comparing

[10] We also test a number of other MDID approaches which yield similar results that are available on request.

[11] The "propensity score paradox" refers to the increase in imbalance, inefficiency and bias caused by PSM's attempt to approximate a completely randomised experiment as opposed to a fully blocked randomised experiment (King & Nielsen, 2016).

changes across a single time period it is difficult to make a strong judgement either one way or the other. With that in mind, we proceed to consider DID results for the full sample but we do so tentatively as the possibility of a violation of parallel trends cannot be ruled out.

Tests of balance in the pre-treatment period, 2013, are presented in Table 2. We find quite convincing balance in the full sample, there are some ways in which treatment and control clearly differ but these values are largely either statistically insignificant or very small in magnitude. There are two exceptions to this in internet and change in SW share. While this in and of itself is easily controlled for by their inclusion in our specifications, the indication it provides that perhaps treatment and control also differ in unobservable ways is more worrisome. The 1 provider or less and 1 provider samples show very poor balance and so we don't proceed with their inclusion in any analysis as we find the evidence supporting the violation of our assumptions to be too convincing in their case. The RD designs also invoke some concerns over balance though to a slightly lesser degree; there exist large and/or highly significant differences between treatment and control in terms of income, change in income and education. Other than include them as controls there is little one can do to resolve this problem and so when considering the results, of the RD designs in particular, one must be wary that our validity assumption may be violated.

We find that the our instruments are very relevant in the main sample with a Wald F statistic of 97.9 and we see promising evidence that the equation is not underidentified (See Table 3). We see similar results for the three RD designs with the only difference being a significant decrease in the strength of the instrument; the Wald F statistics still remain around 20 and our instruments are still shown to be highly significantly correlated with internet. Evidence in favour of the null hypothesis that the equation is not overidentified and that the instruments are therefore valid is slightly less convincing. While we consistently fail to reject the null, the p-values are scantly above 0.1 and, when combined with the poor balance of the RD designs, we worry about the veracity of our validity assumption in their case in particular.

Table 2: Descriptive Statistics and Pre-Treatment Balance

	Full Sample		1 Provider or Less		1 Provider		RD		RD (Donut)		RD (Covered)	
	Mean	Balance	Mean	Balance	Mean	Balance	Mean	Balance	Mean	Balance	Mean	Balance
Internet	0.14	0.05***	0.01	0.00	0.01	0.00	0.04	0.00	0.04	0.01	0.05	0.00
Income	740	15	464	-152***	469	-197***	559	-91***	561	-119***	602	-62*
Consumption	585	26	403	-29	415	-16	464	-18	461	-48	487	-12
SW share	0.14	0.00	0.13	0.00	0.14	0.00	0.13	0.01	0.13	0.01*	0.12	0.00
Estrato	1.23	0.08***	1.00	0.00	1.00	0.00	1.07	0.00	1.08	-0.01	1.09	0.00
Rural	0.59	0.01	1.00	0.00	1.00	0.00	0.87	0.01	0.85	0.02	0.84	0.01
Dirt floor	0.72	-0.05***	0.93	0.09***	0.90	0.12***	0.85	0.01	0.84	0.03	0.81	0.00
Household size	4.44	-0.16*	4.65	0.09	4.59	0.14	4.43	-0.13	4.42	0.14	4.36	-0.17
Age	48.19	0.41	50.06	-0.09	49.46	-1.41	49.18	0.87	49.22	1.62*	48.72	0.94
Female	0.25	0.01	0.18	-0.05	0.19	-0.02	0.20	-0.01	0.20	0.01	0.21	0.00
3º Education	0.06	-0.01	0.02	-0.02	0.01	-0.01	0.03	-0.02**	0.03	-0.03***	0.03	-0.02**
Saver	0.15	0.02	0.14	0.11***	0.16	0.13***	0.15	0.02	0.14	0.01	0.14	0.00
Time Trend												
ΔInternet	0.02	0.03*	-0.04	0.04*	-0.02	0.01	-0.02	-0.01	-0.02	0.00	-0.01	-0.01
ΔIncome	-1	43	-96	-10	-166	-31	-109	67	-84	91	-111	80
ΔConsumption	71	1	39	-9	55	9	59	32	57	16	63	34
ΔSW share	-0.06	0.02**	-0.08	0.01	-0.07	0.02	-0.07	0.01	-0.07	0.02*	-0.07	0.02*
Households	6546		1976		836		2672		2033		1870	

*** $p<0.01$, ** $p<0.05$, * $p<0.1$

[a] All values are for the final pre-treatment period, 2013. The mean value is that of the whole sample and the balance column displays the difference in means between those covered by Tigo and those who are not. A positive value for balance is indicative of a larger mean value of the variable for those covered by Tigo when compared to those who are not; conditioned on being covered at all and region.

[b] Time trend values consider first differences between 2010 and 2013.

[c] Dirt floor is a wealth proxy and saver is a dummy for whether or not the head of household considers themself to be a saver.

[d] Monetary values are presented in 2016 USD adjusted for PPP.

Table 3: IV First Stage & Instrument Strength

	Full Sample	RD	RD (Donut)	RD (Covered)
Tigo x 2016	0.158***	0.237***	0.260***	0.228***
	(0.0117)	(0.0365)	(0.0422)	(0.0384)
Tigo x 2016 x d	-0.0469***	-0.0593***	-0.0720***	-0.0614***
	(0.00342)	(0.00946)	(0.0117)	(0.00993)
R^2	0.044	0.033	0.040	0.036
Households	8,002	3,128	2,410	2,239
Wald F statistic	97.9	21.7	19.9	19.6
Partial R^2	0.020	0.020	0.025	0.024
Underidentified	0.00	0.00	0.00	0.00
Overidentified	0.14	0.18	0.20	0.11

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

[a] The Wald F statistic reported is the Kleibergen-Paap Wald rk F statistic and the corresponding critical value for the Stock-Yogo weak identification test of 10% maximal size distortion is 19.9 (Stock & Yogo, 2005).

[b] Underidentified and Overidentified present p-values for the Anderson LM test of underidentifcation and Sargan-Hansen test of overidentification respectively. In the case of the Anderson LM test the null hypothesis is that the equation is underidentified and in the case of the Sargan-Hansen test the null hypothesis is that the instruments are valid and that they are correctly excluded from the estimated equation.

Another assumption we make when interpreting our LATE estimator is SUTVA and while we are not overly concerned with this assumption as its violation would lead to a downward bias on the magnitude of our estimates, further research might consider the existence of spillover effects. One might hypothesise that agents may share in the information they ascertain from the internet with neighbours, relatives or friends which would indeed violate SUTVA. One way one might test this hypothesis of spillover effects would be to see if distance from a household with internet is predictive of our outcomes conditioning on distance to a Tigo tower, coverage by a Tigo tower and coverage.

A further concern that is worth pointing out before any inference is made is that of measurement error. The values for both of our outcomes and a number of our covariates are self-reported values that require the survey respondent to recall events that took place over the timespan of a month. This is something we have no control over and all we can do is assume that any bias associated with either forgetfulness or posturing infront of the enumerator is consistent for each household across time.

5 Results

5.1 Full Sample

Table 4 presents the results for our analysis of the full sample, displaying OLS, DID and IV estimators for each of our two outcomes. Before discussing our two main hypotheses it is worth noting, with any potential concerns over balance in mind, that the coefficients for being covered by Tigo are consistently small and insignificant for each of our specifications and outcomes. While this isn't conclusive evidence, it might suggest that conditioning on our covariates has accounted for any disparities between treatment and control.

In terms of our first hypothesis, that the internet provides some information i^δ and i^r, such that $\frac{d|\delta-\delta^B|}{di^\delta} < 0$ and $\frac{d|r-r^B|}{di^r} < 0$, we don't find conclusive evidence of an effect one way or the other. Our IV, LATE, estimator is large and negative at -11.5% however it isn't statistically different from 0. Two things should be highlighted with respect to the magnitude of the point estimate. Firstly, it remains unclear as to whether its lack of statistical significance is due to our lack of power or that the true effect really isn't different from 0, especially with the sharpened q-value of just 0.118 in mind. Secondly, and somewhat contradictorily, such a large estimate may somewhat heighten our concerns of a violation of our exclusion restriction. Both of these concerns point towards the need for further research to clarify the nature of this relationship. Another explanation may be that the effects are canceling each other out in the way described when we discussed the differences between the self-employed and salaried workers. Future research might disentangle such a result by segregating the analysis by type of occupation. In contrast, our DID, ITT, estimator is quite small and statistically insignificant which isn't hugely surprising given that the take up of Free Basics in our treatment group isn't likely to be very high. The result indicates that Free Basics seemingly did not have a significant effect on the treated group as a whole.

Our second hypothesis, that $f(\theta) > 0$, seems to be quite strongly supported in the data. Our IV estimate suggests that the use of Free Basics is associated with a 12.2 percentage point increase in the 'share-worthy' share of consumption. This would represent almost a doubling of the mean pre-treatment share and so we could potentially be identifying a quite large effect. While this is a striking result, we can't rule out the possibility that this value is inflated somewhat due to a violation of our exclusion restriction. All three specifications show poor fits which points to the fact that 'share-worthy' share might be very difficult to predict or at the very least that our controls don't do a very good job. We see a marginally significant ITT estimate of 0.5 percentage points which can be seen as further support of the IV result.

5.2 Robustness

Three things must be kept in mind when considering the fuzzy RD results displayed in Table 5. Firstly, and most importantly, there are serious concerns over balance in all three samples and so the invalidity of our instruments is a very real possibility in this case. Secondly, the instruments are significantly weaker than in the full sample and with Wald F-statistics comparable to the Stock-Yogo critical value required to reject the null of weak instruments at the 10% level there may be weak instrument bias. Finally, the power is much reduced with samples less than half the size of the full sample. All this culminates in our extreme caution in interpreting these results. With that said the total consumption results are, in terms of their statistical insignificance, wholly consistent with the full sample results. The SW share results are also statistically insignificant but are very small in magnitude which is somewhat contradictory to the full sample result. It is also worth noting that the RD samples also seem to be much noisier with standard errors almost double those of the full sample at around 20% and 5 percentage points for the consumption and SW share outcomes respectively. Due to the statistical concerns outlined one cannot read too much into these results, all we can say is that they don't seem particularly supportive of either our hypotheses or our headline results.

In the case of the matched difference-in-differences we see small and statistically insignificant coefficients for both outcomes. Again, this isn't surprising given the modest take up of treatment but it does cast some doubt on our headline ITT estimate of a 0.5 percentage point increase in the 'share-worthy' share.

Table 4: Full Sample Results

Outcome	Log Consumption			SW Share		
Estimator	OLS	Dif-in-Dif	IV	OLS	Dif-in-Dif	IV
Internet	0.0475***		-0.115	0.012***		0.122***
	(0.0132)		(0.0920)	(0.003)		(0.0229)
			[0.118]			[0.001]
Tigo x 2016		0.0227			0.00530*	
		(0.0145)			(0.00306)	
		[0.133]			[0.133]	
Covered	0.113	0.109	0.118	-0.022	-0.0224	-0.0246
	(0.0862)	(0.0867)	(0.0885)	(0.017)	(0.0166)	(0.0179)
Covered by Tigo	0.0191	0.0135	0.0152	0.015	0.0140	0.0179
	(0.0798)	(0.0804)	(0.0816)	(0.016)	(0.0157)	(0.0172)
Log Income	0.199***	0.200***	0.201***	0.001	0.00165	0.000213
	(0.0107)	(0.0107)	(0.0108)	(0.001)	(0.00122)	(0.00130)
Rural	-0.148***	-0.152***	-0.148***	-0.003	-0.00425	-0.00286
	(0.0450)	(0.0453)	(0.0449)	(0.008)	(0.00838)	(0.00882)
Household size	0.0699***	0.0701***	0.0715***	0.003***	0.00343***	0.00224***
	(0.00361)	(0.00362)	(0.00379)	(0.001)	(0.000745)	(0.000804)
Age	-0.00154	-0.00164	-0.00182*	-0.000**	-0.000489**	-0.000276
	(0.00105)	(0.00105)	(0.00107)	(0.000)	(0.000208)	(0.000223)
Female	-0.0741***	-0.0754***	-0.0785***	0.001	0.000824	0.00412
	(0.0221)	(0.0221)	(0.0224)	(0.005)	(0.00490)	(0.00528)
Saver	0.0225*	0.0234*	0.0261**	0.019***	0.0192***	0.0165***
	(0.0128)	(0.0128)	(0.0130)	(0.003)	(0.00282)	(0.00298)
R^2	0.204	0.203	0.195	0.015	0.014	0.020
Households	8,002	8,002	8,002	8,002	8,002	8,002

Robust standard errors in parentheses and sharpened q-values, corrected for the testing of two hypotheses, in square brackets (Benjamini, Kreiger & Yakutieli, 2006).

*** p<0.01, ** p<0.05, * p<0.1

[a] Controls not included in the table: estrato, tertiary education, dirt floor and time dummies.

Table 5: Fuzzy RD & Matched Difference-in-Differences Results

Outcome		Log Consumption				SW Share		
Estimator	RD	RD (Donut)	RD (Covered)	MDID	RD	RD (Donut)	RD (Covered)	MDID
Internet	0.214	-0.0143	0.0848		-0.0156	0.0506	0.00747	
	(0.203)	(0.202)	(0.208)		(0.0481)	(0.0495)	(0.0504)	
Tigo x 2016				0.0241				-0.00296
				(0.0257)				(0.00668)
Covered	0.109	0.0841		-0.0123	-0.0362*	-0.0117		-0.0327
	(0.0926)	(0.140)		(0.151)	(0.0219)	(0.0323)		(0.0224)
Covered by Tigo	-0.0486	0.0961	0.0804		0.0408	0.0555	0.0528	
	(0.115)	(0.0997)	(0.0848)		(0.0276)	(0.0259)	(0.0192)	
Log Income	0.175***	0.175***	0.179***	0.219***	-0.000945	-0.000554	-0.00157	-0.000608
	(0.0130)	(0.0152)	(0.0163)	(0.0185)	(0.00179)	(0.00208)	(0.00233)	(0.00211)
Rural	0.0181			-0.0115	0.000377			-0.00978
	(0.0984)			(0.0757)	(0.0220)			(0.0183)
Household size	0.0760***	0.0822***	0.0717***	0.0608***	0.00515***	0.00509***	0.00505***	0.00463***
	(0.00577)	(0.00680)	(0.00689)	(0.00537)	(0.00110)	(0.00130)	(0.00136)	(0.00115)
Age	-0.000899	-0.00174	-0.00136	-0.00178	-0.000193	-0.000137	-7.87e-06	-0.000333
	(0.00171)	(0.00200)	(0.00210)	(0.00147)	(0.000311)	(0.000361)	(0.000394)	(0.000270)
Female	-0.0332	-0.0475	-0.0208	-0.0849**	-0.00450	-0.00201	-0.00469	0.00811
	(0.0386)	(0.0454)	(0.0450)	(0.0350)	(0.00806)	(0.00898)	(0.0103)	(0.00905)
Saver	0.0335*	0.0438*	0.0300	0.00698	0.0172***	0.0163***	0.0179***	0.0203***
	(0.0200)	(0.0227)	(0.0244)	(0.0214)	(0.00465)	(0.00538)	(0.00577)	(0.00522)
R^2	0.288	0.305	0.271	0.213	0.025	0.026	0.022	0.014
Households	3,128	2,410	2,239	5,026	3,128	2,410	2,239	5,026

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

[a] Controls not included in the table: estrato, tertiary education, dirt floor and time dummies.

6 Conclusions and Implications

This paper set out to test two hypotheses; that internet access changes one's propensity to consume by correcting misinformed beliefs about future payoffs and that internet access increases conspicuous consumption through social media's catalytic effect on our concern over our relative status. To that end we have found causal evidence of a potentially large effect with regards to the latter but we fail to find convincing statistical evidence of the former. While we don't want to oversell our headline result in light of our robustness check's failure to corroborate it, a 12.2 percentage point increase in the share of conspicuous consumption in the consumption bundle would represent, if confirmed in future research, a massive shift in behaviour. Our robustness checks were plagued with issues of statistical validity and so this lack of corroboration should be seen as merely that and not as contrary evidence.

It is unclear whether or not these are favourable results from a policy perspective. Our insignificant ITT estimates suggest negligible effects at the aggregate level which might suggest programs like Free Basics don't really make a difference to the population at large; however, our LATE results suggest that a scaled-up version of Free Basics might have a significant impact on the way the poor allocate their resources. How an increase in conspicuous consumption might interact with other outcomes such as borrowing, saving and income both in the short and long term will determine whether or not this is net positive or negative for social welfare. Some basic accounting would suggest that, given the limited resources of the population in question, more spending on items such as grooming and entertainment probably doesn't represent the optimal allocation of resources from a long-term perspective. If this proves true then it would bring into question the prudence of supporting these kinds of programs, especially in light of their implications for net neutrality.

One must remember that our results pertain to a very specific type of internet access, the limited selection of websites made available though Free Basics, and one might posit a number of ways in which access to the full uncensored internet may interact with our proposed mechanisms. A drastic increase in the choice set of websites might induce users to face the perils of 'overchoice' such that they end up consuming less belief-correcting information (Iyengar & Lepper, 2000). An increase in available websites may also crowd out time spent on social media therefore potentially dampening any effect on conspicuous consumption. Both of these hypotheses remain open research questions.

This paper has contributed to the literature by undertaking the first rigorous analysis of a causal relationship between internet access and the consumption decisions of the poor but there remain numerous avenues through which our analysis might be extended. For reasons previously discussed, one could isolate self-employed and salaried workers, introduce more degrees of freedom regarding the determinants of the cut-off distance of coverage, consider potential spillover effects or consider a myriad of other related outcomes. We find evidence that access seems to encourage excess in the context of the introduction of Free Basics in Colombia, what remains to be seen is whether or not this finding can be confirmed in future work.

7 References

American Tower. (2014). *Introduction to the Tower Industry and American Tower.* Boston, Massachusetts: American Tower.

Avantel. (2018). avantel.co. Bogotá D.C., Colombia.

Banerjee, A., Cole, S., Duflo, E., & Linden, L. (2007). Remedying Education: Evidence from Two Randomized Experiments in India. *The Quarterly Journal of Economics*, 122(3), 1235-1264.

Barreca, A., Lindo, J., & Waddell, G. (2015). Heaping-Induced Bias in Regression-Discontinuity Designs. *Economic Inquiry*, 54(1), 268-293.

Barrera-Osorio, F., & Linden, L. (2009). The Use and Misuse of Computers in Education: Evidence from a Randomized Experiment in Colombia. *Impact Evaluation Series No. 29*. Washington D.C.: World Bank.

Benjamini, Y., Krieger, A., & Yekutieli, D. (2006). Adaptive linear step-up procedures that control the false discovery rate. *Biometrika*, 93(3), 491-507.

Beuermann, D., Cristia, J., Cueto, S., Malamud, O., & Cruz-Aguayo, Y. (2015). One Laptop per Child at Home: Short-Term Impacts from a Randomized Experiment in Peru. *American Economic Journal: Applied Economics*, 7(2), 53-80.

Claro. (2018). Tienda. tienda.claro.com.co. Bogotá D.C., Colombia.

Cristia, J., Czerwonko, A., & Garofalo, P. (2014). Does Technology in Schools Affect Repetition, Dropout and Enrollment? Evidence from Peru. *Journal of Applied Economics*, 17, 89-112.

Czernich, N., Falck, O., Kretschmer, T., & Woessmann, L. (2011). Broadband Infrastructure and Economic Growth. *The Economic Journal*, 121(552), 505-532.

DANE. (2017). *Pobreza Monetaria y Multidimensional en Colombia 2016.* Bogotá D.C.: Dirección Nacional de Estadística.

DANE. (2018). *Estratificación socioeconómica para servicios públicos domiciliarios.* Bogotá D.C.: Dirección Nacional de Estadística.

ETB. (2018). Tienda. etb.co. Bogotá D.C., Colombia.

Facebook. (2015). *Internet.org App Launches in Colombia.* Internet.org. Palo Alto, California.

Facebook. (2018). Internet.org. Palo Alto, California.

Friederici, N., Ojanpera, S., & Graham, M. (2017). The Impact of Connectivity in Africa: Grand Visions and the Mirage of Inclusive Digital Development. *Electronic Journal of Information Systems in Developing Countries*, (Forthcoming).

Galperin, H., & Viecens, M. (2017). Connected for Development? Theory and evidence about the impact of Internet technologies on poverty alleviation. *Development Policy Review*, 35(3), 315-336.

Glewwe, P. & Muralidharan, K., (2016). Improving Education Outcomes in Developing Countries, *Handbook of the Economics of Education.* Elsevier.

GSMA. (2016). *Digital Inclusion and Mobile Sector Taxation in Colombia.* London: Deloitte.

GSMA. (2017). *Country overview: Colombia.* London: GSMA.

Iyengar, S., & Lepper, M. (2000). When choice is demotivating: Can one desire too much of a good thing? *Journal of Personality and Social Psychology*, 79, 995-

1006

Jensen, R. (2007). The Digital Provide: Information (Technology), Market Performance, and Welfare in the South Indian Fisheries Sector. *The Quarterly Journal of Economics*, 122(3), 879-924.

King, G., & Nielsen, R. (2016). Why Propensity Scores Should Not Be Used for Matching (Working Paper).

Kremer, I., Mansour, Y., & Perry, M. (2014). Implementing the "Wisdom of the Crowd". *Journal of Political Economy*, 122(5), 988-1012.

Lohmann, S. (2015). Information technologies and subjective well-being: does the Internet raise material aspirations?. *Oxford Economic Papers*, 67(3), 740-759.

Lotta. (2017). *100 MB to GB - What Can You Do with 100 MB of Data?*. usmobile.com.

Ministerio de Tecnologías de la Información y las Comunicaciones de Colombia (MinTIC). (2018). *Operadores Completa*. Bogotá D.C., Colombia.

Movistar. (2018). Tienda Movil. movistar.co. Bogotá D.C., Colombia.

Rosenthal, R. (1979). The file drawer problem and tolerance for null results. *Psychological Bulletin*, 86(3), 638-641.

Stock, J. & Yogo, M. (2005). Testing for Weak Instruments in Linear IV Regression. In *Identification and Inference for Econometric Models: A Festschrift in Honor of Thomas J. Rothenberg*. Cambridge, UK: Cambridge University Press.

Tigo. (2018). tigo.com.co. Bogotá D.C., Colombia.

Universidad de los Andes, el Centro de Estudios sobre Desarrollo Económico (CEDE). (2018). *Encuesta Longitudinal Colombiana de la Universidad de los Andes – ELCA*. Bogotá D.C., Colombia.

World Bank. (2016). *Exploring the Relationship Between Broadband and Economic Growth*. Washington D.C.

World Bank. (2017). data.worldbank.org. Washington D.C.

8 Appendix

Table 6: A selection of the services provided through the Free Basics platfrom in Colombia (full list available on request)

Service Name	Category	Use
1doc3	Health	A service where doctors provide answers to health-related enquiries.
AccuWeather	Weather	Provides both short-term and long-term weather forecasts.
Agronet	Government & Education	A government website providing comprehensive information on all aspects of agriculture.
Facebook	Social Media	A less data intensive version of the standard Facebook platform.
Messenger	Communication	A medium for free online communication of text, pictures & files.
Mitula	Classified Advertisements	A platform for the advertising of jobs, cars and real estate.
Su Dinero	Finance	Website providing financial information and advice aimed at encouraging saving.
Wikipedia	Education	Open source encylopedia.

[a] Source: (Facebook, 2015)

Table 7: Handset and mobile internet costs by provider

Provider	Handset		Monthly			Weekly		
	COP	USD	COP	USD	Data	COP	USD	Data
Avantel	140,000	108	20,000	16	2GB	5,000	4	500MB
Claro	69,900	54	24,000	19	0.48GB	6,000	5	120MB
ETB	251,910	195	20,000	16	1.92GB	5,000	4	480MB
Movistar	49,900	39	19,900	15	0.6GB	6,500	5	175MB
Tigo	180,000	140	24,000	19	0.48GB	6,000	5	120MB

[a] The handests selected are the cheapest available that are 3G compatible and the data products selected are the cheapest available by the provider that meet the minimum requirement; 100MB per week (Avantel, 2018; Claro, 2018; ETB, 2018; Movistar, 2018; Tigo, 2018).

[b] USD values are 2016 Dollars in PPP terms.