The Adoption of M-Pesa: A Percolation Approach to Network Goods

Lisa Reed Janet Stefanov Zerrin Vural

24 July 2020

Abstract: In 2007, Kenya's mobile network operator Safaricom launched M-Pesa, a mobile phone-based money transfer service. Today over 95% of Kenyan households use M-Pesa, making Kenya one of the first developing countries to fully embrace mobile payment systems. M-Pesa merits further academic investigation due to Kenya's resulting economic growth and reduction of poverty since its inauguration. Here we reference percolation theory from statistical physics to develop a theoretical model of the spread of M-Pesa from 2007 to 2014. We consider M-Pesa a network good that spreads primarily via word of mouth and assume its chance of adoption is determined by the utility a person can derive from it. This utility increases primarily with the number of M-Pesa users in one's social network. We simulate the spread of M-Pesa throughout Kenya by using social network models and measure the goodness of fit of the model. Our model may be useful in analyzing the potential for the propagation of mobile money in other developing countries. We hope our findings will highlight the positive impact to be made by mobile money systems and motivate others to realize similar effects in developing countries.

1 Introduction

1.1 M-Pesa: Mobile Money

To use M-Pesa, an individual must find an M-Pesa agent to deposit or withdraw money from their account. Money transfers are then sent by SMS messages. The recipient need not be a M-Pesa user, though there is a separate and costlier pricing schedule for non-users. As soon as the transaction is finished, the sender receives a SMS confirmation, and the recipient receives a notification about the transfer of money. Currently, there are 37 million active users and 400,000 agents operating in 7 countries: the Democratic Republic of Congo, Egypt, Ghana, Kenya, Lesotho, Mozambique and Tanzania. M-Pesa is most ubiquitous in Kenya, its premier country, and is being used in over 90% of Kenyan households. The prevalence of M-Pesa in Kenya has prompted reviews from economists and scholars.

1.2 Motivation

Mobile money networks, buoyed by the prevalence of mobile phones, have created an opportunity for a new system of financial capital. One success story is M-Pesa in Kenya.[1] M-Pesa is a mobile money transfer service launched in March 2007 that enables users to transfer money electronically through a series of text messages. The introduction of this app facilitated simple money transfers, and allowed for unbanked populations to have access to formal financial resources. By 2016 M-Pesa's widespread use was apparent, with 96% of Kenyan households using the service [13]. Further significance of M-Pesa lies in its impact of having increased per capita consumption and reduced the number of households in extreme poverty [13], lending itself to be a product worthy of further analysis.

Kenya in the period from 2007 to 2014 is an attractive candidate for a percolation model as this time period captures the primary spread of M-Pesa. In this paper, we draw upon statistical physics to model the initial spread of M-Pesa throughout Kenya as a percolation process. Looking further, this percolation model could potentially be applied to predict success or failure in other countries by changing relevant parameters to the data of a given country.

2 Literature Review

Literature on network goods is expansive. There are a few areas that most closely relate to the topics discussed in the percolation approach: microeconomic studies on the household effects of mobile money and econophysics studies on the diffusion of goods.

2.1 Household Effects of Mobile Money

In a seminal paper analyzing the effects of M-Pesa in Kenya, economists William Jack and Tavneet Suri found that mobile money has allowed increased efficiency of both the allocation of labor and the allocation of consumption over time. They estimated that this effect raised at least 194,000 households out of extreme poverty [13]. The ability for a mobile money transfer service to lift 2% of Kenyan households out of poverty is a substantial result that makes further analysis of M-Pesa meaningful and worth pursuing. A better understanding of the dynamics of M-Pesa, such as the method of M-Pesa's spread, could lead to important insights regarding the implementation of mobile money in other developing countries in an attempt to produce a similar result. Studies have found that M-Pesa's resounding success in Kenya is facilitated by the combination of a few factors.

First, Safaricom, a subset of Vodafone and Kenya's main mobile phone provider, has a near-monopoly on telecommunications. Estimates on the market share of Safaricom fall at 88%. This means that Safaricom has a strong initial market presence and is able to efficiently advertise and build participation. Thus, M-Pesa was created on the shoulders of an existing phone network. As a result of the explosion in mobile phone usage, Safaricom agents were already positioned to respond to demand for cell phones. Given the positive correlation between phone use and M-Pesa use, demand for M-Pesa largely mirrors the demand for phone services, and the existing Safaricom network was easily able to scale its operations. Now there are more than 110,000 M-Pesa agents within Kenya, and the average user reports that they can find an agent within 260 meters [14]. Kenyans also report

[1] Pesa is the Swahili word for money.

that M-Pesa is relatively easy to use; on a scale from 1-5, where 5 is extremely difficult, the average rating is 1.479 [14].

Furthermore, public trust in mobile network operators such as Safaricom is higher than public trust in banks and formal financial institutions, consistent with general trends in the developing world. In Kenya, many banks collapsed in the 1990s, creating permanent distrust in the formal financial sector. Another factor contributing to public trust in M-Pesa is its origin as the idea of M-Pesa first started at Moi University in Kenya. By originating within Kenya, it helped create trust as being an app created both by and for Kenyans. Additionally, because nearly all phone users in Kenya are clients of Safaricom as discussed earlier, advertisement through social learning can be very effective and allow M-Pesa to easily expand [7].

Also, Kenyans have found utility in M-Pesa through its facilitation of easy and secure sending of remittances. Studies on the household effects of M-Pesa have found that an urban-rural channel of remittance flows has been supported by M-Pesa [3, 11]. Migrant workers in Kenyan cities using M-Pesa have an efficient and safe method of transferring money to their relatives in rural areas, increasing shared prosperity. This is compatible with the fact that the number of migrant workers in Kenya has risen since the expansion of M-Pesa use [13, 11].

Lastly, Kenya has a substantial unbanked population. For the unbanked, a mobile money service can be viewed as a sufficient substitute for more formal financial services, while also bearing fewer costs and barriers to entry. While M-Pesa is a money transfer service, many users report using it as a means to save by keeping a balance on their account [11].

When considering the potential spread of M-Pesa beyond Kenya, one must keep in mind these circumstances as complicating factors in extending the model to other countries. Regardless, M-Pesa's availability to unbanked populations give it the potential for major success in the developing world. Access to credit and payment systems through this virtual format results in a new sector of the population being able to access financial capital.

2.2 Econophysics Modeling of Diffusion of Goods

It is cheapest and simplest for M-Pesa users to send and receive money within the user base of the app. This feature makes M-Pesa a classic network good: the value of M-Pesa to a consumer grows as more people use it. In the case of M-Pesa, word of mouth advertising is particularly effective and can be considered the primary motivation for which individuals eventually adopt the app.

Previous work by economist Arthur Campbell (2013) models the effects of word of mouth advertisement in the context of different social structures. In particular, Campbell assumes that consumers must receive information about a monopolist's product directly from a user that is also a contact; our analysis continues these assumptions [4].

3 Estimating the Adoption of M-Pesa

In addition to contributing to literature on mobile money networks, we also contribute to literature on network goods generally. We contribute by building a model that can be applied to a network good that relies on widespread adoption and provides valuable insights for targeted advertising strategies.

Specifically, here we summarize data to determine the reasons of adoption or non-adoption of M-Pesa as stated by Kenyan users or non-users.

Our approach to modeling the the adoption of M-Pesa is heavily influenced by the work of Dan Björkegren [2]. In particular, we model adoption as an optimal stopping problem, where individuals delay adoption by some time t, and then they adopt. This technique is appropriate, as adopting M-Pesa has nontrivial barriers to entry for individuals who do not have mobile phones and little incentive for individuals with limited social networks and contacts. We capture both of these effects in our theoretical model.

3.1 Optimal Stopping Problem

In this section, we set up an optimal stopping problem, where the individual decides the optimal time to register as a user of M-Pesa. The individual faces costs of usage, including per-transaction costs, hassle-costs of finding an agent and waiting in line, and potential technology costs such as purchasing a compatible phone. However, the individual derives benefit from performing money transfers. The individual thus earns utility according to the following model:

$$u_{ijkt} = max[1/\phi_{cost} F(n_{i,j,t}, n_{k,i,t}) - \delta(n_{i,j,t} + n_{k,i,t}), 0].$$

In this model, ϕ_{cost} is a constant that incorporates the fixed cost of adoption and a conversion between utility and currency, $n_{i,j,t}$ is the number of transactions from individual i to j in regime t, $n_{k,i,t}$ refers to the number of transactions from individual k to i in regime t, and δ refers to the aggregated per-transaction cost for sending and receiving money, including hassle costs. The function F takes as inputs the number of transactions of individual i and outputs a value for i's utility. The adoption point is the first regime where u_{ijkt} is strictly greater than zero.

With this optimal stopping framework for the adoption decision, we make a few assumptions about the function F. First, F is increasing in both inputs. Second, F in our model is implemented as a single-variable function based on the proportion of individual i's contacts that use M-Pesa. This single variable captures the relative size of the potential circle of candidates for money transfers, and we find it a sufficient proxy for the number of transfers, while also allowing the adoption dynamics and decision to be endogenously determined in each period.

3.2 Motives for Adoption

In reality, there exists never-users, or individuals who will never adopt M-Pesa. This includes both individuals who do not have access to the type of technology necessary to use M-Pesa and individuals that do not have any contacts outside of a small spatial radius and do not travel, and therefore do not have an incentive or need to use mobile money. Never-users are entirely contained in the set of non-users. In our model, we do not include those without access to necessary communication technology as part of the social network. Those with little incentive to adopt mobile money are represented in the network and are reflected in the asymptotic behavior of the percentage of users over time. [2]

[2] This population is limited, given the increasing mobility of the developing world, and it is very difficult to obtain economic activity data for this type of individual.

To determine a baseline estimate of the never-users to compare our results to, we use a pooled estimate of the reasons non-users provided for delaying adoption (Table 1). Since our model assumes a relatively constant proportion of never-users throughout time, we minimized variation over time by using all years to generate an estimate for the upper bound of the percentage of users in the population. We obtained our data regarding non-use and adoption of M-Pesa from the 2014-2016 Financial Inclusion Insights Kenyan survey [9].

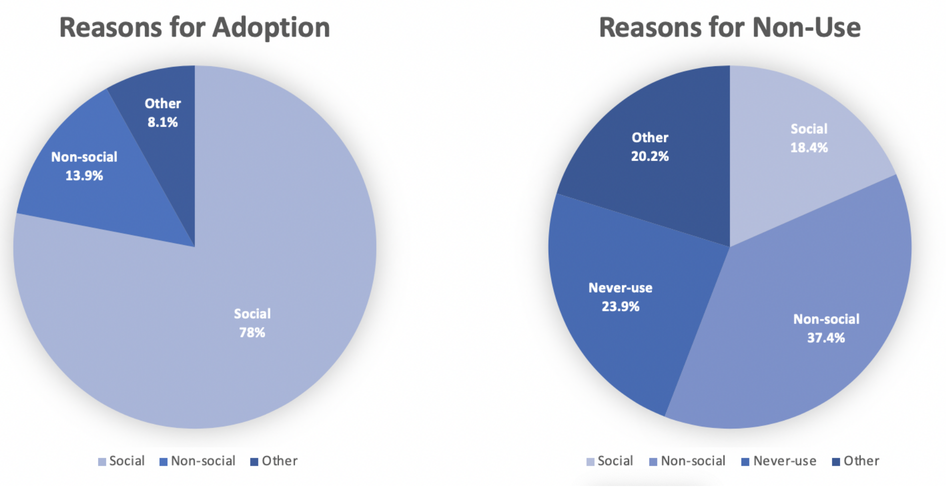

Figure 1: Results from the Insights 2014-2016 Kenya Survey. This survey asked mobile money users the reason for their adoption of it, and non-users the reason for their non-use.

Table 2 shows the reasons users gave for adoption. Since we consider M-Pesa to be spread by word of mouth, we are particularly interested in the proportion of new-users who adopted the app as a result of another user's request. It is notable that 78% of M-Pesa users gave a reason for adoption directly related to interactions with other people, labeled in Figure 1 as social reasons. In contrast, only 18% of non-users gave a social reason for their non-use. This indicates that social factors are a strong determinant of the decision to adopt M-Pesa, but are less influential towards an individual's decision to not adopt M-Pesa. Due to this finding, we are able to justify our assumption that contacts can only positively influence the adoption decision in order to simplify our model. This means that having a contact that is a user shortens the time to adoption. Under full diffusion, the limit set of non-users would be the set of never-users. In the case of partial diffusion, there will be non-users who are not aware that the app is available or who do not have enough social incentive to adopt the app.

Reason	Number	Proportion
I do not know what it is / what I can use it for	52	0.0352700
I do not know how to open one / use the service	130	0.0838170
I do not have a state ID or other required documents	216	0.1392650
There is no point-of-service / agent close to where I live	39	0.0251450
I do not need one, I do not make any transactions	226	0.1457120
Registration paperwork is too complicated	6	0.0038680
Using such account is difficult	29	0.0186980
Fees for using this service are too high	9	0.0058030
I never have money to make transactions with this service	295	0.1902000
No one among my friends or family use this service	8	0.0005158
I do not understand this service; I do not know what I can use it for	30	0.0193420
I do not have a mobile phone	155	0.0999360
I do not trust that my money is safe on a mobile money account	20	0.0128950
My husband, family, do not approve of me having an account	13	0.0083820
It does not provide any advantage over my current financial service	9	0.0058030
Other	314	0.2024500

Table 1: Reasons for Non-Use of M-Pesa

Reason	Number	Proportion
I had to send money to another person	1798	0.248608
I had to receive money from another person	3343	0.465340
Somebody requested I open an account	104	0.014477
I had to send money to an organization/government agency	29	0.004037
I had to receive money from an organization/government agency	18	0.002506
An organization/government agency requested I sign up for an account	20	0.006264
An agent or sales person convinced me	45	0.002784
I saw posters/billboards/radio/TV advertising that convinced me	31	0.004315
A person I know, who uses mobile money, recommended I use it	127	0.017678
I saw other people using it / most of my friends or family are using it	202	0.028118
I wanted to start saving money with a mobile money account	446	0.062082
I wanted a safe place to store my money	443	0.061665
I was given a promotion to start using it	8	0.001114
Other	582	0.081013

Table 2: Reasons for Adoption of M-Pesa

4 Theoretical Model

4.1 A Brief Introduction of Percolation Theory

Percolation theory is a theory of statistical physics that describes phase transitions and the critical behavior of a networked interactions. Things like the internet, power grids, conductivity in materials, porous substances, and virus spread in a population are all examples of systems whose behavior relies on the interconnectedness of its components.

A standard model of percolation theory can be considered with the spread of a forest fire. Given a lattice, each node is occupied by a tree with probability p, and is not occupied by a tree with probability $1 - p$. The fire spreads if there are trees occupying the sites next to a tree on fire. This eventually reaches a fixed state where the spreading stops as no nodes catch fire. Increasing p results in clusters of trees on fire. There exists some critical probability P_c needed in order to go from local clusters to global behavior. This is an example of site percolation.

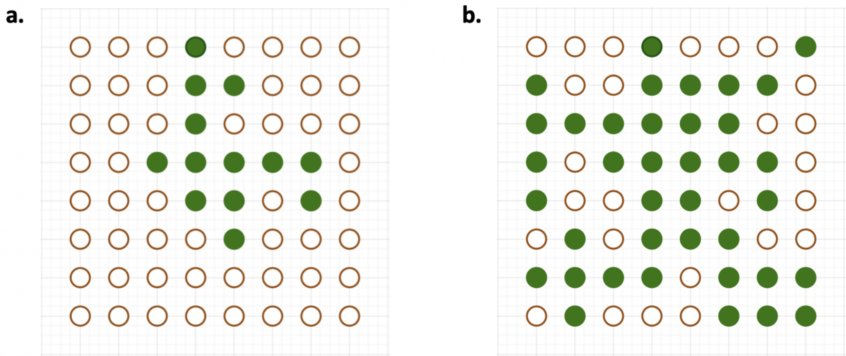

Figure 2: Site percolation modeled by trees (green circles) occupying each lattice point of a forest with probability p. p can also be though of as the packing ratio of the $2D$ lattice. In panel a, $p < P_c$, and therefore any fire that may break would be contained. In panel b, $p > P_c$, so a forest fire would percolate from the top of the lattice to the bottom, spreading throughout the forest.

Percolation theory has also been used to model the spread of disease through communities as illustrated in Figure 3. In this models, a node represents a host for a disease, and is occupied if the host is susceptible to the disease. Edges between nodes represent possible contacts of disease transmissions, and the edges exist between hosts with probability p. If the hosts are also susceptible with some probability (i.e. some nodes will not be occupied, thus slowing the spread of disease), then this is an example of site and bond percolation. With this, the critical behavior of a disease going from being contained to becoming an epidemic can be studied.

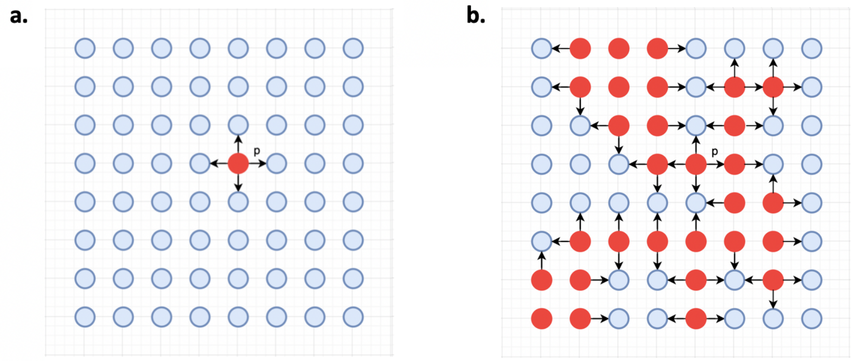

Figure 3: Here bond percolation is illustrated by a simple $2D$ lattice. Let every person in a lattice be susceptible, and assume they are in contact with the nodes adjacent to them (nodes are not connected diagonally). Red nodes represent infected people. Let p be the probability that an infected person transmits the disease with someone they're in contact with. In panel a, $p << P_c$, one person is infected, and has not spread the disease. In panel b, $p > P_c$, and now several people have become infected. The potential of infected people to spread the disease further is represented by the arrows to blue, uninfected nodes.

Percolation theory has also been used in economics to model the spread of a consumer good. In these cases, the good is usually given a utility function that factors into whether or not someone adopts a good. We study the percolation of M-Pesa across Kenya as a network good, since the value of M-Pesa grows as more people use it.

4.2 Population Networks

The interconnectedness of Kenya's population is represented by a network where nodes are people and edges are connections between people. Since we are looking at the demand formed through social learning of consumers, it is important that this network has properties of a real world social network. A number of different structures were considered to model a social network, notably Erdős–Rényi, Watts-Strogatz, and Barabási-Albert networks.

Different types of networks have different types of topology, such as the nodal degree, which refers to how many connections each node has. Regular random networks are statistically homogeneous in the degree of nodes, and in the pattern of connectivity of nodes. However, most real world networks don't follow the homogeneous nodal or connectivity distributions found in regular and random networks. Erdős–Rényi (ER) random graphs feature a set of discrete nodes that are connected by edges with uniform and independent probability. These types of network structures produce nodes with statistically homogeneous nodal degree, and the degree distribution is Gaussian. These graphs tend to have short path lengths and low levels of clustering. This is not consistent with real world social networks, which are not generated randomly, typically have larger clustering coefficients, and have degree distributions that follow a power law - and therefore do not have a homogeneous number of connections amongst the nodes.

Networks that exhibit what is called the "small world property" have a small average path length, meaning the number of edges on the shortest path between nodes is small compared to what you would see in a regular or random graph. Sociologist and mathematician Watts and Strogatz wanted to more accurately model real world social networks, which were found to have high clustering and short path lengths and created a generative model for this. The Watts-Strogatz network structure attempts to create a "small-world" effect by connecting discrete nodes with non-uniform, pre-selected probabilities and then randomly rewiring them. This creates a high degree of clustering, leading to groups with high probabilities of connectivity among individuals. The Watts-Strogatz structure also allows for "long distance" connections resulting in short average path lengths, which is consistent with social patterns. This structure has attractive properties, but it does not capture the organized variability in the connectivity of individuals as seen in real world networks.

Finally, the Barabási-Albert structure gave the most realistic and desirable properties in the context of the model. This graph generates scale free networks by using preferential attachment, meaning new nodes are more likely to be connected to existing nodes with high nodal degrees. As a result, a small but significant number of nodes have a very large degree (called hubs), while a large number of nodes have small degree. Evidence supporting this "rich get richer" phenomenon in social networks can be found in the data and microdata from Kenya. Furthermore, the Barabási-Albert network's scale free property means that its degree distributions follows a power law, where the probability of a node having a given degree is $P(n) \sim n^{-\gamma}$ where $\gamma > 1$. The scale-free characteristic of this network is more consistent with real world networks, and allows us to create a network that represents both village and urban populations. Furthermore, in previous work epidemiologists have made use of the Barabási and Albert model to capture social interactions amongst susceptible populations. This supported our decision to adopt this type of network structure for social connections [5, 1, 10].

4.3 The Model

For our percolation model, we consider a population of size N as a network with $N = \{1, ..., n\}$ nodes, connected using preferential attachment to form a Barabási-Albert network with some initial nodal degree k_0. Each person is assigned an initial utility value for M-Pesa according to a normal distribution $u_{i_0} \sim N(\mu, \sigma^2)$. This utility is akin to the probability that someone will adopt M-Pesa if introduced to it. If a person's valuation for M-Pesa is above a certain threshold U_c, then we say that they will adopt the app. Initially there will be a small fraction of the population that adopts the app due to product advertisement. The rest of the population will learn about the app through word of mouth advertisement from their contacts. At each time stamp probabilities were increased by a function dependent on how many people with M-Pesa someone is now connected to, using the formula $u_t = u_{t-1}$ (1 + friend weight $\times$ number of friends with M-Pesa). If the updated probability is greater than U_c, that person will then adopt M-Pesa. [3]

4.4 Model Parameters

In determining parameters, the following terms were defined. Threshold Utility (U_c) is the minimum utility someone must receive to adopt M-Pesa. This can also be interpreted as the total costs of use for the intended number of transactions. Initial Utility, (u_i), is the initial valuation of M-Pesa for a

[3] The software that we used to generate the random graphs and compute the adoption probabilities at each stage can be provided upon request.

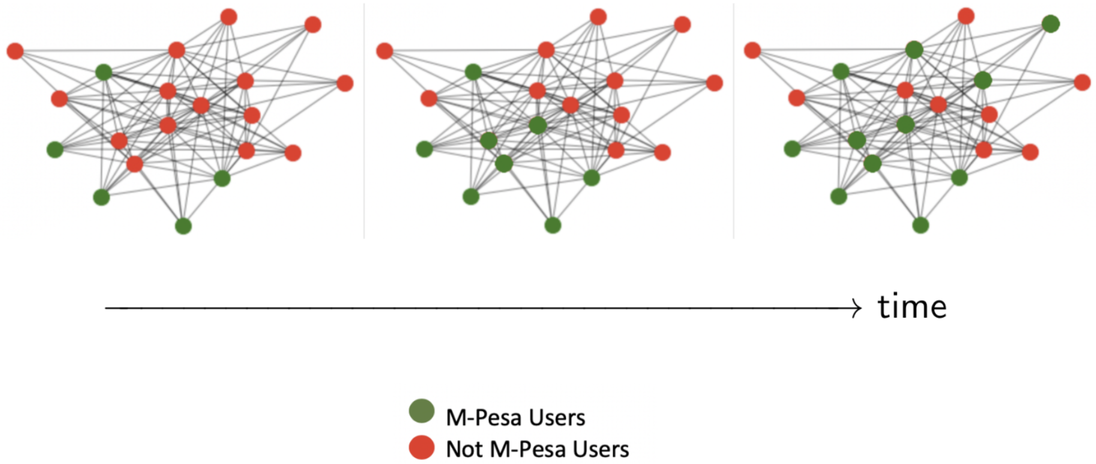

Figure 4: At each time stamp we expect to see an increase of M-Pesa users especially from individuals whose contacts have adopted M-Pesa in recent time stamps, consequently increasing the utility that person can derive from adopting M-Pesa, making them more likely to adopt. In this diagram we expect red nodes who's number of green conacts are increasing over time to have a higher probability of adopting M-Pesa in subsequent months.

user, which determines the baseline level. We used a normal distribution for these initial utilities. An adjacency matrix represented the connections between individuals. A one for the entry (i,j) in the matrix shows that individuals i and j are connections. Because most social connections are reciprocal, we assume the connections are reciprocal and thus the adjacency matrix is symmetric. K represents the number of connections in the social network. The friend weight determines a user's percentage increase in utility for each connected friend that adopts M-Pesa.

Our model depends on the individual's utility changing at each stage in response to the decisions of their contacts in the previous stage. The impact of having friends who are users of M-Pesa on the individual adoption utility follows a pre-selected function that remains constant throughout the model, but the inputs are specific to each individual. In creating our initial constraints, we assume monotonic increases in utility. A linear change was selected in the adoption utility based upon each additional connection that adopts M-Pesa.

5 Implementation & Results

In developing code for the percolation simulation, we adapted sections from Allen Downey's book *Think Complexity* [6] to build the social network. The initial parameters to set up in the code are as follows: the number of nodes, average number of edges, initial probability mean and standard deviation, threshold number U_c, friend weight, and length of run time in months.

Data from Jack and Suri's M-Pesa round 5 survey was used as a baseline estimate of the spread of M-Pesa [14]. We summed the number of new adopters and cumulative percentage monthly from March 2007 through July 2014. Figure 5 shows the percent of adopters through the time period. Looking at the data, it seems to follow a logarithmic curve.

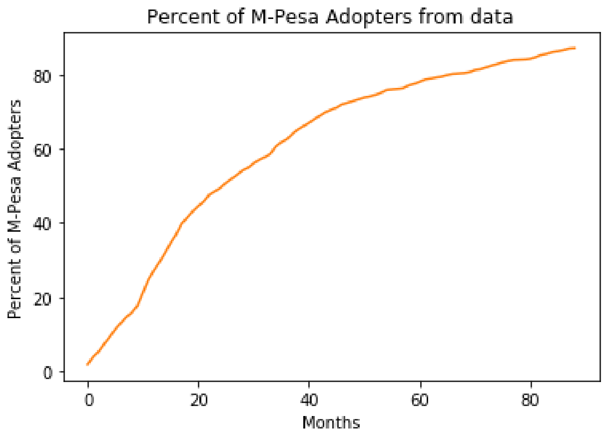

Figure 5: Using data collected from Jack and Suri, we plotted the increase in the percent of M-Pesa users out of Kenya's total population over 89 months.

Multiple simulations were run with the results plotted below. Each simulation was plotted against the actual data to better show its fit.

Our first run created a strongly exponential growth, quickly reaching an asymptote at 100% as shown in Figure 6 panel a. This did not fit the data well, as the output was more of a logistic curve.

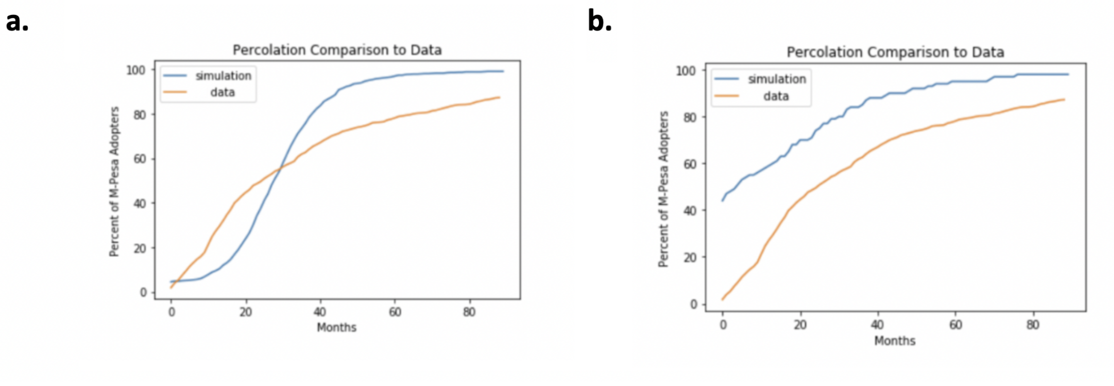

Figure 6: Panel a) Trial 1 Panel b) Trial 2

After the first trial, our goal was to eliminate the initial exponential behavior of our model, and to try to fit the concave down portion of the trial to the empirical data's logistic curve. To do this we decided to run a simulation ignoring the original parameters. The dataset reports approximately a 2% adoption rate in the first month, therefore the initial constraints had been set to achieve a similarly low percentage for initial adopters. By ignoring this value, and letting the starting percentage be around 44%, a curve was obtained that followed closely about 20 percentage points above the empirical data curve as shown in Figure 6 panel b.

Thus far, the best fit run has been with $1,000$ nodes, 25 average edges, friend weight of 0.001, initial utilities distributed with mean 0.21 and standard deviation 0.1, and a threshold utility of 0.3.

Best Fitting Parameters	
Population Size	1,000
Initial Nodal Degree	25
Friend Weight	0.001
Initial Utility Distribution Mean	0.21
Initial Utility Distribution Standard Deviation	0.1
Threshold Utility	0.3

The resulting curve from inputting these parameters can be seen in Figure 7. A low threshold value is to be expected, given the low costs of using M-Pesa compared to the benefits of money transfer.

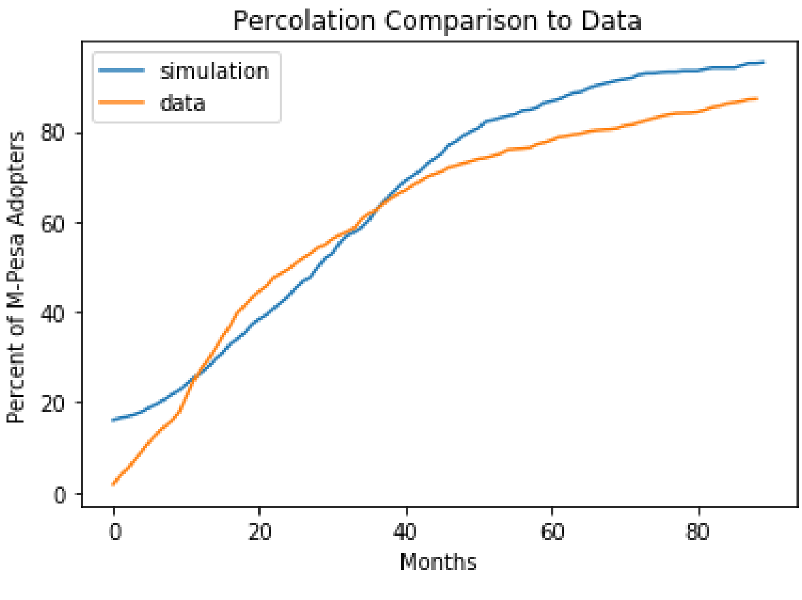

Figure 7: Trial 3

5.1 Conclusion

We performed a Kolmogorov-Smirnoff bootstrap test on the difference between the data from the Jack and Suri dataset and our simulation results. The software for the bootstrap came from the Matching Package in R developed by Sekhon [12]. Running the test initially, starting at January of 2008 because our data on the initial nine months of M-Pesa adoption is somewhat incomplete, the point estimate of the difference in the two distributions was 0.42105 (p-value $2.815 * 10^{-6}$).

Upon closer examination of the results, we recognized that there were two potential components of distance between the distributions, namely a vertical difference and a shape difference. Because we are ultimately interested in modelling how M-Pesa spreads, we uniformly shifted our simulation's adoption percentages vector by 9 percent, or the amount of the vertical difference at the end points. When we ran the Kolmogorov-Smirnoff bootstrap on this new distribution, the point estimate for the difference in the distributions to be 0.14474 (p-value 0.4036). This change in the results of the nonparametric test indicate that our model effectively captures the pattern of M-Pesa adoption within Kenya. Furthermore, evidence from the World Bank's MENA Economic Monitor suggests that 96% of Kenyan households have an M-Pesa account [8]. Our model reaches full diffusion at 96.3% , while the Jack and Suri data indicates full diffusion at 87.2% of the population, though the Jack and Suri dataset ends at 2015, while the World Bank's Data is from 2019.

Additionally we were able to estimate the proportion of never users to be 23.8% of non-users by adding together the non-use reasons given in the tables. Under the full diffusion assumption the percentage adoption of only non-users should be equivalent to the percentage of never users. In the 2017 FII index data, 87.2% of individuals surveyed were users of M-Pesa. Combining these two pieces of information, our point estimate of the size of the never-user population is 3.02% (the proportion of the survey's non users that are never users). This indicated to us that the full diffusion model should approach a horizontal asymptote at 97% country-wide adoption. If we continue to run our simulation past the time frame of 89 months, the model asymptotically approaches an adoption level of 97.5%. The latest data on M-Pesa in Kenya show that adoption exceeds 96% of households using M-Pesa, which suggests that our model and method of capturing the never-users and the resulting steady state behavior is consistent with data. The next round of survey data from the Brookings Institution's Financial Inclusion Index will provide further information about the diffusion of M-Pesa.

6 Discussion

We have considered several ways to extend our study further. Exploration of different function that govern changes in a person's utility of adoption may yield more accurate results. Ideally, we would like to have regional demographic data to better assign utilities based on each individual's livelihood, social group, and personal factors to better tailor our model to the specific population studied. Also of interest is the relationship between number of nodes, and average number of edges. To most accurately model the spread, a reliable estimate of the number of close connections each individual has is necessary. Our estimate was based on anecdotal data, but a more complete and robust estimate would yield results that are more accurate with respect to the initial network setup. We would also like to include in our model the occurrence of adding (and subtracting) new contacts over time, as would happen in a more realistic social network.

In general, the dynamics of updating utility values in response to the changes in an individual's contacts determine the long-run behavior of M-Pesa adoption more than changes in initial conditions and parameters. This is an important consideration for firms looking to bring network goods to the market. We believe that mobile money networks in the developed world should be examined through the lens of a diffusion model, and this is a direction for future research. Our model of diffusion is not limited to M-Pesa and is useful for any network good. However, the assumptions that govern the dynamics of adoption for M-Pesa will likely change in accordance with properties specific to the good being diffused. M-Pesa is unique in that registering as a user of the app carries relatively low cost, in addition to the lower price schedule for users. This unique set of characteristics mean that an individual is unlikely to discontinue their use of the app in a formal way. Our monotonicity assumption relies on the pricing and system of M-Pesa, but other network goods may not follow this general pattern.

Further studies of real world social network structures could be used to optimize targeted advertising toward individuals on the part of the firm. The ideal candidate for advertising based on the network structure is another topic of interest, with potential applications of machine learning techniques.

Additionally, M-Pesa has been brought to other markets in Africa, Latin America, East Asia, and Eastern Europe. It was largely unsuccessful commercially. If data can be collected to assign individuals' utility value based on individuals' demographic information instead of randomly assigning them based on a normal distribution, the diffusion model can then be applied to analyze why full diffusion did not occur in specific countries. Similarly, this technique could be applied to prospective markets to predict the extent of diffusion once mobile money is introduced.

References

[1] Shweta Bansal, Bryan T Grenfell, and Lauren Ancel Meyers. When individual behaviour matters: homogeneous and network models in epidemiology. *J R Soc Interface*, 4(16):879–891, 2007.

[2] Daniel Björkegren. The adoption of network goods: Evidence from the spread of mobile phones in rwanda. *The Review of Economic Studies*, 86:1033–1060, 2019.

[3] Judith Bosibori Bosire. M-pesa: Why kenya? Master's thesis, State University of New York at Buffalo, 2010.

[4] Arthur Campbell. Word-of-mouth communication and percolation in social networks. *American Economic Review*, 103(6):2466–2498, 2013.

[5] Leon Danon, Thomas A. House, Jonathan M. Read, and Matt J. Keeling. Social encounter networks: collective properties and disease transmission. *J R Soc Interface*, 9(76):2826–2833, 2012.

[6] Allen Downey. *Think Complexity, version 2.6.3.* 2016.

[7] Lauren Dunn. What leads to a mobile banking program's success?: A comparison of m-pesa and eko india financial services. *Journal of Public and International Affairs*, pages 108–117, 2015.

[8] Caroline Cerruti François de Soyres, Mohamed Abdel Jelil and Leah Kiwara. What kenya's mobile money success could mean for the arab world. https://www.worldbank.org/en/news/feature/2018/10/03/what-kenya-s-mobile-money-success-could-mean-for-the-arab-world. Accessed: 2020-08-07.

[9] The Financial Inclusion Insights Program Intermedia. Fii kenya datasets, 2013-2017.

[10] Sebastian Meyer and Leonhard Held. Power-law models for infectious disease spread. *The Annals of Applied Statistics*, 8(3):1612–1639, 2014.

[11] Olga Morawczynski and Mark Pickens. *Poor People Using Mobile Financial Services : Observations on Customer Usage and Impact from M-PESA.* CGAP Brief. World Bank, Washington, DC, 2009.

[12] Jasjeet Singh Sekhon. *Multivariate and Propensity Score Matching with Balance Optimization*, 2013. R package version 4.8-1.2.

[13] Tavneet Suri and William Jack. The long-run poverty and gender impacts of mobile money. *Science*, 354(6317):1288–1292, 2016.

[14] Tavneet Suri and William Jack. M-PESA Kenya Rounds 1-5, 2017.