

Mobile Money in Tanzania*

Nicholas Economides[†]

Przemyslaw Jeziorski[‡]

July 31, 2015

Abstract

In developing countries, mobile telecom networks have emerged as major providers of financial services, bypassing the sparse retail networks of traditional banks. We analyze a large individual-level data set of mobile money transactions in Tanzania. Transactions can be classified as (i) money transfers to others; (ii) short distance money self-transportation; and (iii) money storage for short to medium periods of time. We utilize a natural experiment of an unanticipated increase in transaction fees. We find that the demand for long-distance transfers is less elastic than for short-distance transfers, which suggests that mobile networks actively compete with antiquated cash transportation systems in addition to competing with each other. Further, we find that the willingnesses to pay to avoid walking with cash an extra kilometer and to avoid storing money at home for an extra day are 1.1% and 1% of an average transaction, respectively, which demonstrates that m-money ameliorates significant amounts of crime-related risk. We explore the implications of these estimates for pricing and propose Pareto superior price discrimination.

JEL classification: O16, O17, O33, L14, L15

Keywords: mobile money network, financial exclusion, transaction costs, Tanzania, banking, social network, crime

*We thank the Bill and Melinda Gates Foundation for financial support and MIC Tanzania Limited for providing the data. We also thank Paul David, Ben Hermalin, Ganesh Iyer, Jake Kendall and participants of seminars at NYU Stern, Duke University, University of Minnesota and Northwestern University for helpful discussions.

[†]Stern School of Business, NYU; NET Institute; economides@stern.nyu.edu

[‡]Haas School of Business, UC Berkeley; przemekj@haas.berkeley.edu

1 Introduction

Financial exclusion is seen as an important impediment to growth in developing countries (see Schumpeter, 1912; King and Levine, 1993; and Levine, 1997). Exclusion is particularly prevalent in sub-Saharan Africa as there is a significant gap in financial infrastructure compared to other developing regions (see Demirgüç-Kunt and Klapper, 2012). One consequence of this underdevelopment is that many African households lack access to traditional banking services. Filling this gap, the rapidly expanding mobile phone networks introduced mobile money (m-money) wallets, which are attached to the phone number of the customer and provide many functions of traditional bank accounts. Mobile money provides consumers with access to a relatively inexpensive and reliable way of performing financial transactions that can potentially augment money liquidity and ameliorate crime-related risk.

An expanding literature in marketing and economics demonstrates the positive effects of new institutions and technologies on economic efficiency in emerging markets. Aker (2010) demonstrates that mobile phones reduce dispersion of grain prices in Niger, and Jack and Suri, 2013 use consumer expenditure surveys to conclude that adoption of mobile banking contributes to consumption smoothing and do not find an effect of m-money on savings rates. There is also growing literature in marketing studying similar issues¹, particularly Miller and Mobarak (2015), Sudhir and Talukdar (2015), Qian et al. (2014) and Zhang (2015). However, despite very significant adoption of m-money in many developing countries, there is insufficient analysis of a potential impact of m-money on alleviating financial exclusion.

This paper provides evidence on the impact of m-money on alleviating financial exclusion using a unique data set of millions of m-money transactions procured from Tigo,² a leading telecommunication network in Tanzania. We analyze the transaction patterns and classify the uses of m-money into those that improve financial liquidity and those that alleviate the risk of exposure to different types of crime, such as street crime and burglaries. We explore the implications of a natural experiment created by an exogenous and unanticipated increase

¹See the editorial by Narasimhan et al. (2015) and the survey by Sudhir et al. (2015)

²Tigo is the second largest mobile telecom and m-banking network with about 18% market share. According to the sources within the company, Tigo's network consists of slightly more urban users than other networks.

in the transaction fees. We make three contributions.

First, we make a causal inference about the slope of the demand curve for various m-money transactions by applying the discontinuity approach and examining transaction propensity within a short time horizon before and after the fee increase. The increase in total surplus facilitated by m-money is hard to measure directly, thus, the willingness to pay (WTP) provides a useful revealed preference measure of the economic value of m-money transactions compared to the antiquated outside options. We show which dimensions of financial exclusion and risk amelioration are affected. In particular, the WTP for liquidity augmentation informs us about the extent of economic harm caused by low financial liquidity in sub-Saharan Africa, while the WTP for risk amelioration informs us about the economic harm generated by the high level of crime.

Second, we estimate the extent to which consumers' decisions are driven by transaction fees in contrast to non-monetary factors, such as switching costs or limited technological and financial literacy. The relative importance of monetary and non-monetary factors has implications for the design of policies aimed at improving access to financial services.³

Third, after developing a stylized supply model of m-money services, we construct a Pareto superior transaction fees schedule, which can be used by both companies and regulators to improve the efficiency of the m-money system.

Mobile money networks differ from traditional banking, in that each network creates its own m-money that circulates within its network, but cannot circulate outside it. Conversion from cash to m-money ("cashing-in") is free. Transfers across users within the same m-money network are relatively inexpensive, with the fee for the average transfer being 1.1%. However, conversion from m-money to cash ("cashing-out") is relatively expensive, incurring a fee of 7.3% on average.⁴ Further, transfers to users of a different mobile network require the sender to visit a network agent in person, cash-out the transfer amount, later cash-in the same amount in the second network and finally execute the transfer in the second network.

³A motivating setting for this question is examined by Cole et al., 2011, who find that both monetary and non-monetary factors are important for the adoption traditional banking accounts in Indonesia.

⁴Fees are higher as a percentage of small transfers and small cash-outs. The fees reported (1.1% and 7.3%) are the realized averages. If all transfers and cash-outs were made at the average amount of 38,000 Tanzanian Shillings, the fees would be 0.66% and 2.6% respectively.

The total monetary cost of such transaction amounts to 8.4% on average.⁵

In Tanzania, mobile banking has significant adoption, that is, almost thirty-five percent of households have at least one m-money account. Thirty two percent of population use exclusively mobile banking and only 2% have an active traditional bank account.⁶ There are three major m-banking networks: Vodacom (Vodafone) with 53% market share in m-money, Tigo with 18% share, and Airtel with 13% market share. Because making telephone calls and transferring money across mobile telecom (and banking) networks is expensive, many consumers have a different phone appliance or a different SIM card for each phone network that they substitute in the same phone. Additionally, Tigo advertises phone appliances that take multiple SIMs.⁷ Some consumers use m-money for business transactions: 21% of Vodacom M-Pesa users do, as do 12% of users of Tigo & Airtel. Cashing-in and cashing-out of m-money is done through a network of fixed and roaming agents that act as ATM machines. This is the main cost of the m-money networks.

We identify three uses of m-money networks which resolve different dimensions of financial exclusion. The first function is the ability to execute instantaneous peer-to-peer (P2P, person-to-person) transfers, compared to the alternatives of transporting money in person, using a bus driver or using Western Union. P2P transfers are predominantly used to transfer remittances from urban areas to rural locations.⁸ Approximately 30% of users with a Tigo account make at least one transfer a week.

⁵Mobile-banking networks are “incompatible” with each other and with the national currency, but there exists a costly “adapter” that provides compatibility at a price of 7.3%. This is similar to traditional technical incompatibility when adapters are available. So, although the two currencies are incompatible, they become convertible (“compatible” in the networks literature terminology) at a cost. See Economides (1991), Farrell and Saloner (1992), Katz and Shapiro (1994), and Economides (1996) for discussion of adapters in the compatibility decision.

⁶Source: Tanzania – Quicksights Report FII Tracker Survey

<http://finclusion.org/wp-content/uploads/2014/07/FII-Tanzania-Wave-One-Survey-QuickSights-Report.pdf>

⁷There is an evidence that SIM cards are relatively inexpensive. See, The Financial Inclusion Tracker Surveys Project (2013).

⁸According to The Financial Inclusion Tracker Surveys Project (2013) only 14% of all households made or received a non-remittance payment in the past six months using any type of cash delivery, including using m-money. The most common types of non-remittance payments are school fees, government fees and taxes, utility bills, and salaries.

The second function of the m-money network is to securely carry money for short distances, typically up to 10 kilometers. These transactions, which constitute at least 13% of all transactions, involve depositing to the m-money account (cashing-in) using a local agent, and withdrawing the money at another location (cashing-out), without making a P2P transfer. These transactions are aimed at minimizing the risk of being robbed while walking with medium and large amounts of cash. To the best of our knowledge, such transactions have not been identified before in the economics literature. Transactions incur a relatively high average 7.3% fee, which suggests that the risk of walking with cash is substantial and that mobile money mitigates a considerable amount of that risk.

The third function of the network is to store money for short and medium periods of time. Compared to peer-to-peer transfers and transportation transactions, savings transactions are not very prevalent as 90% of the money leaves the network within 5 days of being cashed-in. Moreover, less than 1% of users keep the money in the network for longer than a month, which indicates low contribution of mobile money transactions to long-run savings rates.⁹ We find evidence of heterogeneous usage of the above services. In particular, transporting and storing money is positively correlated with sending a transfer, and negatively correlated with receiving a transfer.

To describe the demand of mobile services, we estimate a structural model of demand for P2P transfers, and for money transportation and storage. We use a panel data set on all m-money transactions from December 2012 to February 2013 from Tigo Telecom. To identify the slope of this demand curve, we employ a discontinuity strategy based on exogenous and unanticipated change in transaction costs in January 2013, and restrict our attention to transactions within two weeks around that change. The change involved raising the transfer and cash-out prices for certain transfer and cash-out transactions, which affected more than 50% of transactions and triggered an immediate substitution effect. In particular, comparing the week before the price increase to that after the price increase, the aggregate share of transfer transactions that experienced a price increase decreased by 0.3%, and the average transfer distance increased by 1%. This data variation identifies the slope of the demand

⁹According to The Financial Inclusion Tracker Surveys Project (2013), only 20% of population believe that m-money can be used in different ways, for example, to save money.

curve and shifts in demand related to the distance span of the P2P transfers. Additionally, we utilize the within-person variation in the propensity to transact before and after the price change to identify a flexible model of customer heterogeneity. Modeling consumer heterogeneity shifts the identification burden from cross-sectional to within-person variation in the data, which produces estimates that are less susceptible to self-selection bias.

We find that, on average, the consumers (senders) of P2P transfers are price-inelastic. However, we also find that these customers are heterogeneous and respond quite differently to changes in fees. In particular, those customers who execute large transactions are usually more price inelastic than consumers who execute smaller transactions. We suspect that this difference is an outcome of the income effect, but can also be an outcome of larger value for liquidity of more wealthy customers. Additionally, we find that demand for long-distance transfers is less elastic than for short-distance transfers. This is consistent with the price of the traditional alternatives, such as using a bus driver, being greater for long-distance transfers than for short-distance ones. Thus, despite significantly lower transaction fees when using any of the m-money networks, we believe that many marginal customers are choosing between antiquated money transfer means and a particular m-money network (in this case Tigo), rather than choosing between two competing m-money networks. Indeed, more than 70% of Tigo users report that they have never used another network (see The Financial Inclusion Tracker Surveys Project, 2013). This softens competition between mobile networks, and suggests that non-price factors play a significant role in the overall adoption rate of m-money.

Consumers executing transportation and storage transactions are, in contrast to the consumers of P2P transactions, moderately price-elastic. This difference may be related to more urban penetration of transportation and storage transactions, which results in better access to substitutes. Using these estimates, we compute the consumers' willingness to pay for risk amelioration achieved using transportation and storage transactions. We find that consumers are willing to pay up to 1.1% of the transaction amount to avoid walking an extra kilometer with cash, and up to 1% to avoid storing cash at home for an extra day. Thus, we provide the first monetary estimates of an economic harm, based on revealed preferences from field data, caused by high levels of street crime and burglaries. These estimates of damages suggest

a significant loss of welfare arising from poor law enforcement in Tanzania and are largely consistent with an extreme level of crime as reported by United Nations Office on Drugs and Crime (2009) and The Financial Inclusion Tracker Surveys Project. Large estimated levels of WTP suggest both that crime risk is high and more importantly that m-money ameliorates this risk to a significant extent. In addition, judging by the relatively high popularity of transportation and storage transactions in urban areas, we suspect that such areas are particularly affected by this type of economic harm and that is where the m-money is the most effective because of the dense agent network.

We find that a sender of a P2P transfer takes into account, to some extent, the cash-out fee paid by the receiver. This suggests that the incompatibility of m-money with other forms of money, including m-money of other phone companies, as evidenced by large cash-out fees, has a negative effect on the propensity to make P2P transfers. We find that the network would realize a higher revenue if it decreased cash-out fees for the receiver and simultaneously increased transfer fees. However, because some users frequently use the network to transport or store money without making a transfer, decreasing the overall cash-out fees would decrease revenue from such transactions, and would lead to an overall decrease in profits of the network. Therefore, without price discrimination, users who transfer money subsidize users who do not transfer. We propose a feasible and incentive-compatible price discrimination scheme that solves this problem. Namely, the network should charge a zero cash-out fee for withdrawals that do not exceed a recently received transfer amount (if any). For all other cash-outs, the network should charge a positive cash-out fee that is slightly smaller than the transfer fee. This scheme, coupled with a simple fixed mark-up pricing of transfers, delivers a Pareto improvement in which both the network's profits and consumer surplus from transfers, transportation and storage are higher than those under the current pricing.

2 Data

The data set contains all mobile financial transactions among subscribers of a major cellular phone service provider in Tanzania for the months of December 2012, and January and February 2013. Each record contains a user ID, balance pre- and post- transaction, type

of transaction and the location identifier. There are six types of transactions: peer-to-peer transfer (P2P), balance check, cash-in, cash-out, bill payment, and recharge of the phone account. We focus our attention on P2P and cash-out transactions.¹⁰ To distinguish among short-, medium- and long-distance transfers, we match the mobile transaction data with GPS data on cellular location IDs in Tanzania obtained from `locationapi.org`.¹¹

As mentioned earlier, P2P transfers and cash-out transactions carry fees. According to anecdotal evidence from Tanzanian sources, changes in these fees are infrequent and unanticipated. In our data, we observe one such change on January 24th, 2013, and we report the tariffs before and after the change in Tables 1 and 2. We utilize these tariffs' change to identify the demand for m-money transactions. We use the full three-month sample to identify the population of m-money users. We define a user as a person who made at least one transaction between December 2012 and February 2013. To minimize the potential price endogeneity issues, we employ a discontinuity strategy and utilize only transactions that occur a week before and after the price change. Namely, we utilize the fact that the price change acts as a quasi-experimental variation under the assumption that the exact timing of the tariff increase was uncorrelated with other time-varying demand factors. Thus, the period immediately before the price change acts as a control group and period immediately after the price change acts as a treatment group. Consequently, the average effect of the price increase informs us about the elasticity of the demand. In addition, we explore the data variation generated by the fact that not all transaction types experienced a price increase. In

¹⁰We drop phone account recharge transactions because they are executed at no charge to customers.

¹¹We are able to match 75% of the cellular IDs into fine-coded GPS positions. The remaining 25% of location IDs are inferred in the following way. For each sender, we obtain the previous active location (within a few hours), and match that location ID to the GPS data. The previous transactions can include any activity excluding being a receiver of the P2P transfer, for which we do not directly observe the location data. In this way we are successful in matching 24% of the data. The remaining 1% is approximated by using the first few numbers of the location ID, which defines the location of sender with an error of a few miles. For our objective to distinguish short-, medium- and long-distance transfers, the matching procedure is reliable. Since we do not observe the location ID of the receiver, we have to approximate it. Using a procedure similar to one we used for the unmatched senders' IDs, we match the most recent transaction of the receiver. In the vast majority of cases, we are able to match the receiver within a few hours of receiving the transfer, usually cashing-out or checking the balance of the account.

particular, we are able to identify a common time trend in m-money transactions separately from the effect of the fee increase using a Method of Moments approach which explores the same data variation as difference in difference regression.

Table 3 reports aggregate statistics about the executed P2P transfers. We utilize a subsample of 1,400,644 unique consumers observed over the period of 14 days. We include customers that are registered but made no transactions over the focal 14 days, which minimizes self-selection problem usually present in data sets that include only executed transactions.¹² The average transfer is approximately 38,000 Tanzania Shillings (TShs), or about 24 USD. The transfers are large compared to the 2007 average monthly consumption of 58,000 TShs.¹³ and average monthly consumption of m-money users of 130,000 TShs.¹⁴ Moreover, the transfers are quite dispersed, with a standard deviation of 84,000 TShs. These statistics suggest that a large portion of a user's household income is channeled through a m-money platform. Figure 1 contains the histogram of all Tigo P2P transfers. We note that the amounts are clustered at the edges of the intervals, which suggests that choosing the transfer amount would be more reasonably approximated by a discrete versus a continuous choice. Further, on an average day, approximately 4% of all people with a Tigo m-money account make one transfer, and approximately 1% make two transfers. The high frequency of transfers coupled with large transfer sizes suggests that part of the household money flow may not be captured by the official consumption figures.

The average geographic distance of transfers is 70km with a standard deviation of 110km. The full distribution of transfer distances is depicted in Figure 2. A large percentage of transfers are short distance; however, we also observe a fat tail of transfers that span hundreds of kilometers. Figure 3 presents a logarithmic frequency of the transfer origination by geographic location. The transfer origination points are concentrated in more developed

¹²Due to a short time-span of our data, we are unable to capture new adoption of Tigo's service that results from price changes. Our estimates of the slope of the demand for a price increase should be unaffected; however, our estimates of the impact of price decrease should be regarded as short-run. We note that short-run demand elasticity is likely to be smaller than long-run demand elasticity. We discuss the consequences of this difference when computing pricing counterfactuals.

¹³See United Republic of Tanzania, Poverty and Human Development Report (2009).

¹⁴See The Financial Inclusion Tracker Surveys Project (2012)

regions, such as Dar es Salaam and Zanzibar Island. We also observe significant transfer activity along roads and rivers. Figure 4 depicts the topology of the first one thousand transfers in the data. The star network topology confirms that Dar es Salaam and Zanzibar are the main sources of the money flow. We find that than 90% of extremely short-distance transfers (less than 10 meters between sender and receiver) occur at the center of the star in either Dar es Salaam or Zanzibar. Moreover, receivers of such transfers carry an average pre-transfer balance of 50,000 TShs, whereas the balance of the receivers that are further than 100 km from the sender averages only 25,000 TShs. Thus, we believe that the short distance transfers are likely related to commercial activity in the urban areas.

In the following section we present the descriptive analysis of the data.

3 Network description and reduced-form analysis

Because we observe all financial transactions within the network, we can track an m-money Shilling from birth (cash-in) to its death (exit from the network). For this purpose, we tag an m-money Shilling cashed-into the network at the beginning date of our data (December 1st, 2012) and simulate its path through the network. We find that only about 1% of Shillings do not exit the network during the 3 months of our data. The money leaves the network relatively quickly, with 60% of Shillings being cashed-out within 2-3 days of being cashed-in (see Figure 5). The money exits the network predominantly in the form of regular cash-outs (about 85% of cases). The other exit ways include: topping up of the mobile telecom account (7%), electricity bill payment (6%), and other in-network purchases such as telecom data bandwidth packages (1%).

Figure 6 shows the distribution of the number of P2P transfers (hops) during the lifetime of a dollar in the Tigo network since cash-in. The majority of money in the network gets transferred only once and then leaves the network, 10% exits after two hops, and the percentage that stays three or more hops is miniscule (see Figure 6). As we noted earlier and depicted in Figure 5, the lifetime of money in the network is short.¹⁵ This suggests that

¹⁵The inflation rate in Tanzania is relatively low and averages to 6% in the relevant period. Thus, we exclude high inflation rate as a possible cause of the low lifetime.

the m-money network is used predominantly for one-time transfers and not as a full-fledged financial system where the money would circulate extensively from user to user. We come back to this finding when discussing pricing experiments and counterfactuals later in the paper.

We further investigate the instances in which the money leaves the network without ever being transferred (cash-in-cash-out or zero hops transactions). Approximately 35% of funds entering the network do not get transferred to anyone but get cashed-out by the original depositor (who has to pay the cash-out fee). The histogram of the transaction amounts for zero hops transactions is presented in Figure 7. We find that this distribution is quite similar to that of peer-to-peer transfers, suggesting that the relative sizes of the two types of transactions are comparable. For zero hop transactions, we find that the median time the money stays in the network (lifetime) is approximately 2.5 days. We present the full histogram of the lifetime of money for these transactions in Figure 8. The money is cashed-out within 3 hours in 17% of the cases, while it stays in the network more than 10 days in 20% of the cases. Thus, we believe that cash-in-cash-out transactions serve two distinct purposes: (i) as a short-distance money transportation vehicle, (ii) as a short- and medium-term savings account. The short lifetime cash-in-cash-out transactions occur primarily within the capital city of Dar es Salaam. According to United Nations Office on Drugs and Crime (2009), nearly 40% of 1,884 surveyed households in Dar es Salaam reported that they were the victims of theft or robbery in 2007. These crimes include a large share of car, bicycle or motorcycle hijackings. Thus, even driving with the monthly paycheck in cash may not be advisable.¹⁶ Using a cash-in-cash-out strategy can serve as viable insurance against losing the money, and is facilitated by the fact that the agents' network in the capital city is dense.

To investigate these transactions further, we compute a median distance between cash-in and cash-out, conditional on the lifetime. The median distances for smaller lifetime values reflect predominantly money transportation while the median distances for a larger lifetime reflect long-term savings. The medians for transactions with a lifetime smaller than a few hours are presented in Figure 9. We observe that the median travel distance between cash-

¹⁶According to another data set on crime, provided by The Financial Inclusion Tracker Surveys Project, 6% of households experienced theft/robbery/burglary or assault in the last 6 months as of 2012.

in and cash-out is increasing as a function of lifetime.¹⁷ In particular, the transactions with extremely short lifetimes of less than 20, 40 and 60 minutes have a median traveling distance of 5.2, 6 and 6.2 kilometers, respectively. In these cases, the user cashes-in, boards a means of transportation and immediately cashes-out on arrival.¹⁸ Considering that the average cash-out fee is approximately 7%, the prevalence of cash-in-cash-out transactions can be rationalized by at least 7% probability of being robbed, which is modest in the light of aforementioned crime statistics. We note that the median distances are constant beyond the life span of 3 hours, which indicates that the required traveling distance can be covered within this time.

Figure 10 depicts the median traveling distance for zero hops transactions with longer lifetimes. In contrast with short lifetime transactions, the median traveling distance of these transactions is decreasing in the cash-in-cash-out time interval. This reflects the fact that many longer lifetime transactions are not executed to physically transport cash, but rather use the network to keep savings. Indeed, zero hops transactions with a lifespan longer than 10 days have a traveling distance of about 1 kilometer, which suggests that users cash-in and cash-out within the same mobile network cell area. Keeping money in the network provides no interest. However, mobile network accounts are generally considered safer than keeping money at home. About 13% of respondents in the UN survey report being a victim of burglary or attempted burglary, which puts in question the long-term safety of keeping savings at home.¹⁹

More information about the structure of the zero hops transactions can be obtained by examining the joint distribution of distance and lifetime in the network, which is depicted in Figure 11. Savings transactions reside in the upper-left part of the figure (short distance, long lifetime), and transportation transactions reside in the upper-right part of the figure

¹⁷ These traveling distances are based on the distance between the center of the cell at cash-in and the center of the cell at cash-out. For small traveling distances it is possible that the user moves from an edge of one cell to an edge of another cell, or moves within the cell.

¹⁸For security reasons agents are obliged to enforce identification checks during cash-out. Thus, cash-out by a person other than the account owner is prohibited.

¹⁹This is also indicated by focus group studies by Plyler et al. (2010) in Kenya. There is considerable evidence that the crime rate is higher in Tanzania.

(long distance, short lifetime). The remainder of transactions occupying the interior of the histogram serve both savings and transportation purposes. The existence of such transactions suggests that savings and transportation should be modeled jointly.

The above discussion suggests that cash-out preceded by a transfer can be considered as a separate product from a cash-out not preceded by a transfer. To investigate this further, we compute the probabilities of cash-out originating from a peer-to-peer transfer or from a cash-in by the same customer.²⁰ As shown in Figure 12, the distribution of funding sources of cash-outs is bimodal, that is, the cashed-out money is either funded almost entirely by cash-ins or peer-to-peer transfers, that is, we observe little mixing of money from different sources on the mobile account. Exploiting this dichotomy, we label cash-outs funded more than 50% by a cash-in as cash-in-cash-out transactions.²¹ Based on the earlier discussion, we distinguish two purposes of such transactions: (i) money transportation, and (ii) savings, and investigate the distance between cash-in and cash-out. In particular, the cash-outs within less than one day of a funding cash-in have a median distance span of $8.5km$, indicating short-distance transporting motives. Similarly, the cash-outs after more than one day of a cash-in have a median distance span of $3km$, indicating a savings motive rather than a money transportation motive.

We find positive correlation between the number of send transfers and the number of transportation and storage transactions (Pearson correlation coefficient of 0.1 with $p < 0.001$). However, we find negative correlation between the number of received transfers and the number of transportation and storage transactions (Pearson correlation coefficient of -0.07 with $p < 0.001$). It indicates that people using the network to transport and store money are also initiating transfers, but the people who receive transfers are less likely to transport and store. These correlations are related to the fact that the senders are usually located in urban areas, where transportation and storage is relatively more popular.

Summing up the above discussion, in the remainder of the paper we examine the demand for two products. First, a transfer product consisting of a cash-in, peer-to-peer transfer

²⁰We compute the probability by simulating backward path of each dollar cashed-out from the network between January 17th and February 3rd, 2013. Using this procedure we can trace back the origin of about 97% of the cashed-out money.

²¹The actual cut-off has negligible impact on the results in the remainder of the paper.

and a cash-out. Second, a transportation-savings product consisting of cash-in followed by cash-out by the same user. The transfer product and the savings product resemble retail banking products in developed countries; however, the transportation product is unique to Tanzania, and possibly other developing countries, and according to our knowledge has not been identified in the literature. Examination of the demand for these products provides comprehensive insights into the role of mobile banking in the growth of developing countries.

We now describe the network pricing specifics. On January 24, 2013, Tigo changed the transfer fees by splitting the transfer amount band 4000-9999 into two bands (4000-4999, and 5000-9999) and doubling the fee for the higher band. It similarly split the band 10000-49999 to 10000-19999 and 20000 to 49999, and increased the fee on the high band by 25%. It also increased cash-out fees in some bands (see Tables 1 and 2).²²

Figure 13 depicts the market shares of each transfer bin before the fee increase and change in the market shares as a result of the fee increase. As shown, the price increase affected most popular transfer bins. For this reason, even if most of the bins are unaffected, more than 50% of transactions executed before the fee increase would have had higher transaction costs after the fee increase. We observe the lower transfer band experienced .15% loss in the demand, while the higher transfer band experienced only 0.2% loss in the demand as a result of the fee increase. Also, we find evidence of substitution to adjacent bins, with bin 50,000-249,999 experiencing relatively large market share increase. Moreover, we obtain regression discontinuity estimates of price elasticities, which amount to $-0.0127(0.0062)$ for the interval 5,000-9,999 and $-0.0177(0.0232)$ for the interval 20,000-49,999. These estimates should be treated as back-of-the-envelope approximations, because they do not take into account switching to outside option and changes in the cash-out fees. Nevertheless, they are suggesting that we do have enough variation in the data to estimate the slope of the demand.

²²The non-linear nature of the pricing schedule offers a cheaper possibility to execute certain transfers (on the left side of the pricing intervals) by splitting one transfer into multiple transfers. Our data does not contain many instances of instantaneous multiple transfers between two individuals; nevertheless, our structural model captures some of the splitting as a substitution to a lower transfer bin, and a possibility of executing more than one transfer in a given period. We account for the fact that possible splitting generates correlation of the error terms within consumers, and across transfers, but we do not account for possible time dependence, which would require a complicated dynamic model.

As an effect of Tigo’s fee increases, customers may switch to using the outside option (typically bus driver carrying cash) or to using rival Vodacom’s m-money network. To develop an insight into the extent of switching into Vodacom, we can explore two parts of the data. For the relevant transfer bounds, which are 4,000-9,999 and 10,000-49,999 TShs, the Vodacom pricing schedule is the same as Tigo’s pricing schedule before the price change. This is useful for our analysis because Tigo’s fee increase amounted to 50 Shillings in both price bins. However, the fee increase, relative to the old fee and the Vodacom fee, amounted to a 100% increase for the 4,000-9,999 TShs band and only a 25% increase for the 10,000-49,999 TShs band. Under the assumption that the m-money networks are not systematically different except for pricing, we should observe much more switching to Vodacom in the lower transfer band. Indeed, the magnitude of switching relative to the bin market share is smaller in the higher bin. Moreover, the price of Vodacom does not depend on the transfer distance, while the price of the true outside option does so. Thus, in the world where there is no switching to the traditional outside option we should observe no change in the average distance of the Tigo transfer after the price increase. In the data, we observe that the transfer distance does change as an effect of the Tigo fee change. In particular, average transfers in the bins where the price has increased span, on average, 1% longer distances.

In order to investigate the descriptive variation in the demand for transfers, we estimate the following reduced form demand equation $q_{zt} = \lambda_0 + \lambda_1 p_{zt} + \epsilon_{zt}$, where $z \in \{1, \dots, 11\}$ indexes transfer brackets from Table 1, t indexes days within the examined two-week window, q_{zt} is a daily number of transfers and p_{zt} is the price of the transfer. We estimate λ_1 to be -5.6 , which is statistically significant with the $p < 0.01$. This coefficient does not change and remains significant after including day dummies.²³ Price coefficient of -5.6 translates to approximately 0.8% decrease in the demand for transfers due to observed price increase. Note that such analysis does not account for substitution across transfer bands treating each transfer band as separate product, and it does not account for consumer heterogeneity. Nev-

²³ When using product dummies together with time dummies we obtain a price coefficient of -2.9 ; however, the coefficient is not statistically significant because the number of aggregate observations is small relative to size of the price effects. When estimating the main model we employ individual-level data and obtain extra power by controlling for consumer heterogeneity.

ertheless, the negative price coefficient provides preliminary evidence about the magnitude of the demand elasticity. In the next section we present a structural model that accounts for substitution and heterogeneity.

4 The Model

In this section, we provide a structural model of peer-to-peer transfer transaction and cash-in-cash-out transactions. We start with a m-money network with N subscribers. Each subscriber needs to transfer money to another person, transport money over some distance, and securely store money as savings.

4.1 Demand for peer-to-peer transfers

Each day t , user n has up to J money transfer needs indexed by j . The transfers include remittances and mobile payments. Each transfer (t, j, n) has a recipient m . Conditional on wishing to make a transfer, the user may execute it using Tigo or transfer the money using an outside option which includes other m-money providers. If the transfer is made using Tigo, we model the decision of how much to transfer as a discrete choice, because, as explained in a Section 2, the amounts transferred are clustered at the edges of round intervals. Consequently, we discretize the size of the transfers into F discrete choices, indexed by f . Each choice is characterized by a transfer amount a_f .

Each transfer is characterized by a the distance to the recipient d_{tj}^n as well as the recipient's current account balance b_t^m . Both of these quantities are assumed to be known to the sender. This assumption is close to reality for transfers between family members. If the transfer is a payment for a service, the sender is likely to correctly estimate that the balance of the receiver is large (and therefore the recipient's average percentage cash-out fees are small). Further, the sender has beliefs about the average daily cash-out strategy of the recipient as a function of the recipient's account balance, b_t^m . Using these beliefs, the sender can compute the size of a cash-out fee imposed on the recipient as the difference between incurred daily cash-out fees before and after the transfer. Formally, the fee imposed on the recipient when

sending a_f Tanzanian Shillings is given by

$$e_t^n(a_f, b_t^m) = \mathbb{E}_t[\text{daily cash-out fees}|b_t^m + a_f] - \mathbb{E}_t[\text{daily cash-out fees}|b_t^m],$$

where both expectations have subscript t to stress that they are conditional on the current cash-out tariff.

The utility of executing a mobile transfer f in the Tigo network is given by

$$u_{tjf}^n = v_f^n + r_t + \alpha^n p_t(a_f) + \beta^n e_t(a_f, b_t^m) + \epsilon_{tjf}^n, \quad (1)$$

where v_f^n is the baseline utility of the transfer f , and r_t is the time trend, $p(a_f)$ is a transfer fee for transferring a_f of m-money expressed in Tanzanian Shillings, and ϵ_{tjf}^n is a random shock to the utility of the transfer. The parameter α^n represents a disutility of the sender associated with paying a transfer fee p_t . The parameter β^n captures a disutility of the sender associated with imposing a cash-out fee e_t on the receiver.

The term r_t captures a daily seasonality effect and an overall time trend in the utility of mobile transfers. Including this term is necessary because we observe that some days of the week enjoy higher transfer rates than others. In particular, Sunday has the smallest transfer propensity, while Wednesday has the highest. Moreover, we need to control for the fact that mobile transfers become more popular over time.

The term v_f^n expresses time-persistent preferences for mobile transfers in general and for transfers of particular sizes. For example, some users may intrinsically prefer to use the outside option, possibly due to lack of an agent network or a preference for another mobile provider. Furthermore, we observe in the data that particular users tend to execute only large transfers, whereas other users tend to execute small transfers. It is important to capture these features, because they impose restrictions on switching patterns.

We model v_f^n using a flexible parametric specification. Let a_f be the amount transferred with a transfer f expressed in thousands of Shillings. We set

$$v_f^n = \bar{v}_f + \bar{\eta}^n + \eta_f^n + \kappa^n \log(a_f). \quad (2)$$

The term \bar{v}_f is a fixed effect capturing the average preferences for particular transfer size f in the population. Random shocks $\bar{\eta}^n$ and η_f^n are assumed to be independent, identically

distributed, mean-zero Gaussian random variables with scale parameters $\sigma_{\bar{\eta}}$ and σ_{η} , respectively. The first random term, $\bar{\eta}$, modifies the utility of all Tigo transfers for a particular user and reflects the persistent preferences for Tigo network, overall. The second random term, η , modifies the utility of a transfer choice f , and reflects intrinsic preferences to send transfers of a particular size. Random coefficient κ^n captures the persistent preferences for sending small versus large transfers. For example, if $\kappa^n < 0$ the user prefers small transfers over large transfers, and vice versa. The introduction of κ^n is motivated by the fact that users who sent large transfers in the past are likely to send another large transfer in the future. The parametric specification for κ^n is described below.

The heterogeneity in the sensitivities to transfer and cash-out fees are captured by the variation in the random coefficients α^n and β^n .²⁴ Mobile transfers involve two fees, the transfer fee and the cash-out fee, thereby adding additional modeling dimension. Some consumers may be insensitive to the cash-out price paid by the receivers and maximize their utility separately from the utility of the receiver. The same users may be still highly responsive to the transfer price. Senders with high price sensitivity in general are likely to have high α^n and β^n . Moreover, it proves crucial to allow users who send large transfers to be less price sensitive. Such possibility is quite natural since the transfer size is likely to be correlated with wealth. To capture these patterns, we model α^n, β^n and κ^n as jointly normal with an estimated variance-covariance matrix, which enables all coefficients to be different, and correlated in a flexible way. We normalize the mean of κ^n to be zero because it is absorbed by the fixed effects \bar{v}_f . We note that when designing the above setup, we tried to write down the most parsimonious model that matches heterogeneity patters in the data and produces sensible counterfactuals. We discuss the further details and goodness of fit in Section 6.

The term ϵ_{tjf}^n represents idiosyncratic shocks to the utility of making a transfer f . In particular, ϵ_{tjf}^n alters users' average preferences for transfers and rationalizes the fact that users do not make the same transfer even if the prices, balance and distance to the receiver are the same. Usually, idiosyncratic shocks are modeled as independent identically distributed

²⁴The importance of the heterogeneity in the price sensitivity in the models with a single price was discussed in Berry et al. (1995).

(IID) normal random variables (probit) or IID extreme-value random variables (logit). Note that ϵ is a residual idiosyncratic shock after accounting for the user-level random coefficients describe above. However, for particular user, the IID assumption would imply that shocks are independent across transfer sizes. This assumption may be too restrictive for the application to monetary transfers. Namely, we observe that when the price of a particular transfer size increases, users substitute disproportionately into adjacent transfer sizes that did not experience price increases. Part of this effect is already captured by the parameter κ . Nevertheless, in response to the increase in the transfer fee and beyond his persistent preferences expressed by κ , a user may decide to send more money than he initially intended in order to reach a lower fee bracket and experience economies of scale. To capture this regularity we allow the adjacent ϵ_{tjf}^n to be correlated. In particular, ϵ_{tjf}^n are distributed as mean zero normal random variables with a tri-diagonal variance-covariance matrix. We place 1 on the diagonal (normalization) and estimate the off-diagonal term ρ . This is a minimal structure that generates correlated preferences for adjacent transfer amounts, which can accommodate disproportional substitutions into neighboring transfer sizes.

The specification for consumer-level coefficients v_f^n , α^n , β^n and ϵ_{tjf}^n can be described as a panel version of a random-coefficient, nested probit model. An observation is a transfer-need/day/user combination. The model is a generalized probit because the adjacent ϵ_{tjf}^n s are allowed to be correlated. Also, it has two levels of nesting: an outside option is a separate nest as a result of allowing for an aggregate shock $\bar{\eta}^n$, and every transfer size is a separate nest because of the persistent shock η_f^n .

The outside option utility includes using competing m-money networks, wire transfer, cash, and asking a bus driver to transport the money. The cost of using competing m-money networks and executing a wire transfer does not depend on distance d_{tj}^m . Wire transfers are very expensive and thus infrequent. The cost of using cash and a bus driver depends on distance. Consequently, we model the utility of the outside option as

$$u_{tj0}^n = u_0 + \gamma d_{tj}^m$$

The fees of competing m-money providers are constant in our sample period and thus absorbed in the constant term. As a consequence, we are unable to identify the cross-price

elasticity of demand across networks.

Because we observe behavior of each user before and after the change in fees, we are able to use differential price responses of users to identify two selection mechanisms. First, we estimate the unconditional distribution of random effects, and thereby are able to separate which types of users are affected by the price change. The crucial identification assumptions are that the knowledge about the fees change is uncorrelated with the type of user, in particular, that the fees' change is unanticipated by all users; and that the knowledge about the fee change diffuses instantly. Second, we estimate an unconditional distribution of the distance between the sender and the receiver. This allows us to assess the selection of executed transfers by distance. We identify this selection using within-user differences in the propensity to execute short- and long-distance transfers before and after price change.

These two selection mechanisms are important for assessing the welfare implications of the changes in transfer fees. Primarily, some users would be less elastic to the fee change, which would result in them internalizing most of the fee increase. Additionally, the demand for long-distance transfers may be less elastic because the outside option is more expensive for longer distances.

We omit the cash-out fees from the utility specification of the outside option. One could argue that once the money is in the mobile account it has to be cashed-out before one can execute the transfer using the outside option. There are two reasons why we think that such considerations are of second order. Primarily, we see the Shillings being cash-out immediately as opposed to accumulated for future transfers. In other words, we rarely see Shillings being sent in the network more than once before they cash-out. Also, most of the peer-to-peer transfers are executed using freshly cashed-in funds as opposed to accumulated ones.

4.2 Demand for transporting and storing money

The model of transportation and storage/savings is similar to the peer-to-peer transfer model. We concentrate our discussion on the differences between the two models. Each consumer n has J needs to transport or store the money, indexed (with a slight abuse of notation) by j . Each transaction is characterized by a distance $d_{S,tj}^n$ and time-span (lifetime in the network) between cash-in and cash-out $s_{S,tj}^n$. The transactions with high distance d and low

lifetime s are money transportation, whereas the transactions with low d and high s are storage transactions. Note that it is possible to combine storage with transportation when both d and s are high. Formally, we model distance and time span as jointly log-normally distributed with an arbitrary variance-covariance matrix.

The utility of executing a transaction of size f in the Tigo network is given by

$$u_{tjf}^{S,n} = v_{S,f}^n + r_t + \beta^{S,n} p^S(f) + \epsilon_{tjf}^{S,n}, \quad (3)$$

where the superscript S distinguishes the parameters of the transportation/storage model from the parameters of the peer-to-peer transaction model. The price sensitivity parameter $\beta^{S,n}$ multiplies the cash-out fee $p^S(f)$. In contrast, β multiplies the cash-out fee in the transfer model. Also, the transfer price is absent because no mobile transfer is executed in the transportation/storage transactions. Other aspects of the inside good utility specification are the same as in the peer-to-peer transfer specification.

The outside option in case of transportation and storage transactions includes moving with money in the pocket either by using public transit, going by foot or private car, storing cash at home, using traditional banking, or using a competing mobile network. The outside option of carrying cash has significant dangers. It is likely that the distance traveled carrying cash affects the probability of being robbed. In a similar way, the number of days for which cash is stored at home affects the probability of losing it through burglary. For these reasons, we propose the following specification for the utility of the outside option:

$$u_{tj0}^{S,n} = u_{S,0} + \gamma^{S,1} \log(d_{tj}^{S,n} + 1) + \gamma^{S,2} \log(s_{tj}^{S,n} + 1),$$

where $\gamma^{S,1}$ and $\gamma^{S,2}$ represent disutilities of walking with cash and storing cash at home, respectively.

5 Estimation

The peer-to-peer transfer model is estimated in two stages. First, we estimate a model of recipient's m cash-out activity at day t as a function of his account balance b_t^m . We postulate the following semi-parametric model for the cash-out fees

$$c_t^m = \theta_0^m + \theta_1 b_t^m + \theta_2 (b_t^m)^2 + \tilde{\epsilon}_t^m, \quad (4)$$

where c_t^m is an incurred cash-out fee at day t by user m . Note that the above equation can be classified as a censored regression model because cash-out fees may be zero if no cash-out actions are executed on day t , that is, c_t^m is censored at 0. In addition to censoring from below, the model is censored from above by the maximum cash-out fee of 5,000. Both bounds are accounted for in the estimation.

Note that a censored regression specification embeds a semi-parametric model of cash-out frequency or $Prob(c_t^m > 0 | b_t^m)$, because it does not assume the distribution $\tilde{\epsilon}_t^m$. In addition, conditional on executing a cash-out, the model predicts cash-out size in a similar way as a linear regression model.

The term θ_0^m represents the baseline propensity to cash-out of user m . If θ_0^m is large, the user prefers cash over m-money and cashes out frequently. If θ_0^m is low, the user tends to keep his money in the mobile wallet. In reality, the cash-out patterns vary from user to user, therefore it is necessary to model heterogeneity in the marginal cash-out propensity θ_0^m .

The term θ_1 represents the marginal effect of the account pre-cash-out balance on the cash-out propensity and intensity. Additionally, the term θ_2 models a second-order effect of the balance on the cash-out. We stipulate that $\theta_1 > 0$, thus, if $\theta_2 = 0$, large balances would be associated with incurring proportionally larger average cash-out fees. However, because the cash-out pricing schedule is concave and carrying larger balances may lead to greater usage of the platform in general, the users with larger balances may incur lower cash-out fees per Shilling. For this reason, the effect of the balance on the cash-out fees is likely to be non-linear, and we predict that $\theta_2 < 0$.

By construction, θ_0^m is correlated with the current balance b_t^m , because clients with large θ_0^m cash-out frequently and tend to carry low balances. We allow for this correlation by modeling θ_0^m as fixed effects. This specification, however, introduces an incidental parameters issue because the model becomes non-linear after accounting for censoring.²⁵ To overcome this problem, we use a GMM estimator developed by Honoré (1992) and extended by Alan et al. (2011) for two-sided truncation.

Next we use equation (4) to obtain estimates of the cash-out fee imposed on receiver m

²⁵We would encounter a similar issue if we modeled cash-out propensity using, for example, a logit model with fixed effects.

by sender n

$$e_t^n(a_f, b_t^m) = \mathbb{E}[\text{daily cash-out fees}|b_t^m + a_f] - \mathbb{E}[\text{daily cash-out fees}|b_t^m], \quad (5)$$

We employ a result from Honoré (2008) that establishes a way to compute marginal effects in the censored regression model with fixed effects. We are interested in the marginal effect $\frac{\partial E[c_t^m|b_t^m]}{\partial b_t^m}$, where the conditional expectation is taken with respect to ϵ keeping constant the fixed effect θ_0^m . We compute the marginal effect of the transfer, which determines cash-out externality in the following way:²⁶

$$e_t^n(a_f, b_t^m) \approx (\theta_1 + 2\theta_2 b_t^m) a_f P(c_t^m > 0). \quad (6)$$

The marginal effect of transferring one Shilling is composed of two parts: (i) a first-order effect $\theta_1 a_f$, capturing the effect of a higher account balance, and (ii) a second order effect $2\theta_2 b_t^m a_f$, capturing lower fees per Shilling (economies of scale) when cashing-out large transfers.

In the second stage, we estimate the demand model using the simulated method of moments following Pakes and Pollard (1989). Specifically, we compute 99 micro-moments from the data and choose the parameters that generate the closest match of the simulated moments. We utilize the following moments:

- Market share of each transfer bin before and after the price change -- 66 moments
- Expected number of transfers at each day of the week -- 13 moments plus one excluded moment
- Mean and standard deviation of the distance for the transfers executed before and after the price change -- 4 moments
- Mean and standard deviation of receivers' balances before and after the price change -- 4 moments
- Probability of sending the same amount of money at least twice during the sample -- 1 moment

²⁶Note that we use $P(c_t^m > 0)$ instead of $P(c_t^m > 0|b_t^m)$, because for small b_t^m the probability of cashing-out is close to zero which produces distortion in the estimates of the externality.

- Interaction between the amount of a transfer and the distance of a transfer -- 1 moment
- Interaction between the amount of a transfer and the balance of a receiver -- 1 moment
- Joint distribution of executing n transfers before the price change and m transfers after the price change, where $n, m \in \{0, 1, 2\}$ -- 9 moments

The moments are simulated 10 times for each customer. We set the maximum amount of daily transfers that the user can execute, denoted as J , to three. Executing more than three transfers is infrequent and constitutes less than 1% of the data.

We discretize the transfer amounts into 33 intervals that follow the modes of the empirical distribution of transfers. In particular, we ring-fence the following intervals: ten intervals every 1,000 TSh between 0 and 10,000 TSh, 8 intervals every 5,000 TSh between 10,000 and 50,000 TSh, 6 intervals every 10,000 Shillings between 50,000 and 100,000 TSh, and 8 intervals every 100,000 TSh between 100,000 and 1,000,000 TSh. We round down the transferred amount in each of the intervals (which coincides with the mode of the interval) and treat the data as discrete.

We employ an identity weighting matrix to obtain initial estimates of the model. Next, we estimate the optimal weighting matrix which is used to obtain the final estimates. Standard errors are clustered on the consumer level.

The transport-storage model is estimated in the same way as the second stage of the transfer model.²⁷ It includes additional moments identifying the joint distribution of transaction distance and lifetime, that is, we additionally match: average lifetime of a transaction before and after the price increase, as well as the average interaction of distance and lifetime for the executed transactions. By construction, the estimation does not have moments concerning the receiver.

6 Results

In this section, we report and discuss the estimates of the transfer and transport-storage models.

²⁷The transport-storage model is estimated separately from the transfer model.

6.1 Peer-to-peer transfer model

Table 4 contains the estimates of the cash-out model. The model is estimated separately, before and after the price changes, to account for differences in conditional distribution of cash-out events caused by a change in the cash-out fees. We find that the coefficient of the recipient’s account balance (in thousands) equals 8.64 before and 8.75 after the fee changes, respectively. In other words, conditional on the cash-out, an extra one thousand Shillings on the receiver’s account increase his expected cash-out fees by approximately 9 Shillings. The difference between the coefficients is statistically significant, which confirms that, indeed, the consumers are incurring larger cash-out fees as a function of balance after the fee changes.

Negative and statistically significant coefficients on the square of the recipient’s account balance suggests that the average cash-out fees are a concave function of the balance. This is evidence of economies of scale in cashing-out as well as of a larger utilization of the balance as transfers or mobile phone top-ups when carrying larger balances. We find that such economies of scale are larger after the price change.

Figure 14 contains a graph of the cash-out fee for a modal transfer of 10,000 Shillings and different pre-transfer balances of the receiver before and after the price change. Note that the fee is a decreasing function of the pre-transfer balance. These results confirm our initial predictions that accounts with larger balances experience a smaller cash-out fee for a transfer of a given size. The difference between the cash-out externalities before and after the fee change is significant. We utilize this variation to separately identify the impact of the changes in transfer and cash-out fees on the propensity to execute a transfer.

Figure 15 reports the transfer size fixed effects. We find that the fixed effects are mostly negative, capturing the fact that the P2P transfers are tail events as measured on a daily basis and that we normalize the utility of the outside option to zero.²⁸ Figure 16 reports the daily seasonality in a transfer activity. We observe that the coefficients are small with the

²⁸As mentioned earlier, we cannot distinguish the situations in which the user does not need to execute a transfer from the situation in which the transfer is executed using an alternative service. Thus, the negative coefficients may indicate that some users, given the current fees, do not need to transfer money every day. We suspect that this concern is less important for smaller and medium transfers, but may be more important for larger and more infrequent transfers.

exception of the coefficient on the Sunday dummy.

The remaining parameters of the demand model for transfers (see equations (1) and (2)) are presented in Table 5. We find substantial heterogeneity in the preferences for Tigo transfers compared to the outside option. That is, the standard deviation of the persistent shock to the utility of any Tigo transfer, σ_A , is estimated at 0.4 as compared to the standard deviation of the idiosyncratic shocks which was normalized at 1. Consequently, a large percentage of the population is averse to mobile banking irrespective of pricing. This is consistent with previous qualitative studies that find non-price barriers to usage of mobile banking, such as insufficient trust and the necessity of technical sophistication of both a sender and a receiver (see Mallat, 2007). Additionally, we find considerable persistence in the types of executed transfers. Specifically, we find that the transfer size parameter κ has a standard deviation of 0.39. An interpretation of this parameter is presented in Figure 17, which depicts the distribution of executed transfers by size for 5 types of consumers, such that $\kappa \in \{-3\sigma_\kappa, -\sigma_\kappa, 0, \sigma_\kappa, 3\sigma_\kappa\}$.

We find that the standard deviation σ_η of the persistent transfer level shocks η_{nf} is low. This suggests that most of the user heterogeneity is related to a transfer size and accommodated by κ .

Covariance between idiosyncratic shocks for adjacent transfers is moderate, that is, ρ_ϵ is estimated to be 0.06. Thus, we find evidence that people substitute to the adjacent transfer bins if the price of the particular bin increases. However, this substitution is somewhat limited, which is consistent with the raw data.

We find that coefficients on the transfer fee and on the cash-out fee are negative and statistically significant. This suggests that both fees have an impact on the propensity to transfer. The negative coefficient on the cash-out fee suggests that the sender takes into consideration, to some degree, the externality he imposes on the receiver when sending him m-money. We find that both price coefficients are highly heterogeneous in the population and highly correlated. In other words, people that are sensitive to cash-out fees are also sensitive to transfer fees. We also find that people transferring larger amounts of money have different price sensitivities. The Pearson correlation coefficient between the transfer price sensitivity α , and κ is equal to 0.5. Also, the Pearson correlation coefficient between the cash-out

price sensitivity β , and κ is equal to 0.53. The conditional means of price coefficients for representative values of κ are presented in Table 6.

We use the estimated price coefficients to compute price elasticities. Because the product mix contains multiple prices, the price elasticity can be computed in multiple ways. Price elasticities that rely on an increase in transfer price of one transfer bin, keeping the transfer prices of other bins constant, are not very informative as to the total demand for Tigo transfers. Instead, we compute elasticities that rely on the impact of uniform percentage change in all the prices on the total demand for mobile transfers. Formally, we compute

$$\frac{\partial \text{demand}(\mathbf{p} + \rho \mathbf{p})}{\partial \rho} \text{demand}^{-1}(\mathbf{p})$$

We compute the elasticities measuring demand: (i) as the probability of a transfer, and (ii) as a total transfer Shilling amount. In the former case, we obtain -0.12 transfer fee elasticity and -0.35 cash-out fee elasticity. In the latter case, we obtain -0.2 and -0.39 , for transfer and cash-out elasticities respectively.²⁹ The above numbers serve as primary evidence that, at the current price point, marginal consumers are more sensitive to cash-out fees than transfer fees.

The estimated elasticities show that the network is pricing on an inelastic part of the demand. Such pricing as well as documented large heterogeneity in the overall value for mobile banking suggests that the network may be engaging in penetration pricing. Another piece of evidence consistent with penetration pricing is a steady upwards price trend in both transfer fees and cash-out fees that we observe during the 2010-2014 time period. We also find that the network prices on the more inelastic part of the demand for transfer fees than the cash-out fees. In other words, the cash-out fees are set to be relatively higher than transfer fees, which may be related to an attempt to lock-in the users' money inside the network. This possibility is discussed in the next section.

The coefficient, γ , representing the contribution of the distance of the transfer to the utility of outside option is negative and statistically significant, which reflects the nature of

²⁹Note that the structural elasticities are obtained using a different formula than those presented in Section 3. Structural elasticities measure total demand response rather than bracket-by-bracket response, which is more relevant for profit maximization. Moreover, the structural estimates of elasticities account for switching to outside option and changes in cash-out fees

the outside option. The longer the distance between the sender and the receiver, the more inferior is the outside option. The result is not surprising considering the fact that the outside option consists of giving cash in person or sending it using a bus driver. Negative impact of the distance on the outside option has implications for price elasticity. Specifically, demand for consumers sending long distance is more inelastic. We present the price elasticities in Table 7. Note that for distances longer than 20km the demand is extremely inelastic. This fact, in conjunction with the anecdotal evidence that the bus driver option is inferior and costs about 10% of the transfer, implies that there is not much substitution to other mobile networks. If substitution to other networks was accounting for the majority of switching, we should then observe a negligible amount of elasticity heterogeneity across transfer distances. Although this result may seem somewhat counterintuitive, it should be interpreted within the cultural reality of Tanzania. Anecdotally, one of the biggest barriers to adoption of m-money is mistrust in the reliability of the operator. Such trust issues may be particularly important for a rural population which is the primary recipient of the long-distance remittances. Once trust with one operator is established, the senders are likely to face switching costs, especially in the short horizon, such as those studied in this paper.

In the next subsection we describe the estimates of the transport-storage model.

6.2 Transport-storage model

Figures 18 and 19 contain the cash-in-cash-out model estimates of the transaction size fixed effects and time trend in the transaction propensity, respectively. The interpretations of these parameters are qualitatively similar to the interpretation of the corresponding parameters of the transfer model discussed above. We find that the baseline utility of the transport-storage action is lower than the transfer action, which is consistent with the fact that the majority of the Shillings in the network are transferred at least once.

Table 8 contains the remaining estimates of the structural model of transport and storage. We find that the users are significantly averse to the cash-out fee, that is, we estimate $\bar{\beta}^S$ to be -0.94 . We find that users are more price sensitive when executing transport-storage transactions compared to the transfer transactions. Consequently, as reported in Table 7, the demand for transport-storage is significantly more elastic than the demand for transfers.

The last three rows of Table 7 contain the estimates of the elasticity of demand for cash-in-cash-out transactions with fixed distance d and lifetime s , isolating elasticity of transporting money from elasticity of storing money. Curiously, we find no significant differences between these two elasticities.

Similar to the transfer model, we find significant consumer heterogeneity in price sensitivity, utility for the transfer size, and correlation between price sensitivity and transfer size. These results are compatible with the earlier estimates of the peer-to-peer transfer model, indicating that both groups of users are similarly heterogeneous.

We find that the covariance of the shocks ϵ for adjacent transactions is somewhat larger in the model for transport and storage than in the model for peer-to-peer transfers. We suspect that in the case of transport-storage the consumer can fully control the amount of the transaction, and therefore is able to decrease the transaction amount when the fee increases.³⁰

We find that both the distance coefficient and lifetime coefficient are significantly negative. This suggests that walking with money or storing money at home carries significant disutility. While it is hard to interpret these numbers in isolation, they suggest that people are willing to pay extra not to walk with cash or store cash at home. We quantify this in the next section.

7 Pricing Counterfactuals

In this section we use the demand model to conduct pricing counterfactuals. We are interested in quantifying the importance of transfer and cash-out fees on the propensity to use the Tigo m-money platform. We group the results in several subsections. First, we evaluate the impact of the observed increase in fees on the demand. Next, we study the impact of abolishing the cash-out fees and introduce price discrimination. In the last subsection, we use the estimates of the cash-in-cash-out model to quantify the consumers' willingness to pay to avoid carrying cash in the pocket and keeping it at home.

³⁰An adjustment would require the user to keep some money in his pocket while walking home, or store extra Shillings in the mattress.

Because we are unable to estimate cross-price elasticities, we keep competitors' prices fixed through this section. This assumption should be taken into account when interpreting the results about the price discrimination in sections 7.2 and 7.3, but should have no impact on our estimates of the WTP in section 7.4.

7.1 Impact of fees increase on demand

We start by evaluating the impact of the observed increase in fees (on January 24th, 2013) on the propensity to conduct peer-to-peer transfers. We recompute the demand for transfers keeping the time fixed effects constant at the level before the price change, and we then impose the fee changes. We present the results in Table 9. We estimate that the observed increase in transfer and cash-out fees lowered the demand by about 3% and the impact of the increase in transfer fees alone lowered the demand by about 2.5%.

7.2 Impact of eliminating cash-out fees (full compatibility)

In the next counterfactual, we isolate the impact of receiver lock-in caused by the cash-out fees on the propensity to transfer using the Tigo network. In particular, we change the cash-out fee of sending P2P transfers and recompute the demand. Figure 20 presents the results when setting the cash-out fee to zero for transfers less than or equal to a certain amount. We observe that abolishing the cash-out fee for small transfers (under 1,000 Shillings) does not have a meaningful impact on the demand for peer-to-peer transfers. Abolishing the cash-out fee for transfers less than 10,000 Shillings results in about 10% increase in demand. Even though people who send these smaller transfers are quite price sensitive, the cash-out fees they pay are quite small as well. Setting cash-out fees of medium transfers to zero has a more dramatic impact on demand. In particular, nullifying the cash-out fee on transfers of less than 50,000 Shillings and 100,000 Shillings amounts to 53% and 73% extra transfer demand, respectively. Lastly, zeroing the externalities for the largest transfers has negligible impact on the demand, because the consumers sending these transfers are not price sensitive. All in all, we find that the probability of executing a transfer using the Tigo network increases by 78% if the cash-out fee is uniformly set to zero.

Another interpretation of this counterfactual is as a proxy for compatibility across m-money providers and between mobile providers and cash. Imposing full compatibility would allow costless flow of funds from m-money to cash, which can be proxied by setting the cash-out fees to zero, keeping in mind that conversion from cash to m-money (cash-in) is already free. The very large impact of setting cash-out fees to zero suggests that the potential for regulation of cash-out fees and imposition of full compatibility. However, the issue becomes more complicated, as we discuss in the next subsection.

7.3 Supply model and price discrimination

Abolishing cash-out fees has significant consequences for the profitability of the network. In particular, it results in per user loss equal to 130% of current per user profits. A significant contributor to the above losses is the expansion of transportation/storage transactions. Indeed, under zero cash-out fees, the customers would make approximately 400% more cash-in-cash-out transactions that generate no revenue but carry significant marginal cost. In consequence, users that who make transfers would subsidize users who do not make transfers. Stated differently, one of the reasons for the network to set high cash-out fees is to recoup the marginal cost incurred in cash-in-cash-out transactions. Thus, the network may want to price discriminate by charging a positive cash-out fee to users who do not transfer.

For the above reasons, the estimated 78% increase in transfer demand found in the previous section is an upper bound of the actual effect of setting cash-out fees to zero. To provide a more robust estimate of the welfare gain from decreasing cash-out fees, we propose the following supply model.

Let \mathbf{p}^T be the vector of the transfer fees. Also, denote as \mathbf{p}^C and $\mathbf{p}^{C,S}$ the cash-out fees charged to the receiver of the transfer and to the user executing a cash-in-cash-out transaction, respectively. Under the current pricing, the two cash-out fees are equal, $\mathbf{p}^C = \mathbf{p}^{C,S}$. To obtain the expression for variable profits from transfers, we make the simplification that all the money is immediately cashed-out after a transfer. As shown in Section 2, under the current cash-out fees, this assumption is approximately satisfied, and is likely to remain satisfied under the counterfactual pricing policies considered in this paper. Per user, daily

variable profits from transfers are equal to

$$\pi(\mathbf{p}^T, \mathbf{p}^C) = 3 \sum_{f=1}^F D_f(\mathbf{p}^T, \mathbf{p}^C)(\mathbf{p}^T(f) + \mathbf{p}^C(f) - MC(f)), \quad (7)$$

where D_f is a probability of executing a peer-to-peer transfer transaction of size f , as described in the model in Section 4. The term $MC(f)$ is the marginal cost of the transfer transaction which is equal to a sum of cash-in and cash-out fees of agents and super-agents, which we observe in the data. The expression is multiplied by 3 because in Section 4 we assume that there are 3 transaction needs per day for every user.

Per user daily variable profits from cash-in-cash-out transactions are

$$\pi^S(\mathbf{p}^{C,S}) = 3 \sum_{f=1}^F D_f^S(\mathbf{p}^{C,S})(\mathbf{p}^{C,S}(f) - MC(f)), \quad (8)$$

where D_f^S is a probability of executing a cash-in-cash-out transaction of size f . Total daily profits are given by $N(\mathbf{p}^T, \mathbf{p}^C, \mathbf{p}^{C,S})[\pi(\mathbf{p}^T, \mathbf{p}^C) + \pi^S(\mathbf{p}^{C,S})]$. The number of users N depends on pricing, since the adoption decision also depends on pricing.

We use the above supply model to study the pricing incentives of the network when it considers setting cash-out fees to zero. First, we consider a situation in which the network can price discriminate between users that receive transfers and those that make no transfers. In this case, the firm keeps the $\mathbf{p}^{C,S}$ at the current level and sets $\mathbf{p}^C = 0$. We consider a simple and easy to implement fixed mark-up pricing policy for transfer fees. Specifically, we set $\mathbf{p}^T = \delta \mathbf{MC}$. We find that, if the firm marks up transfers' marginal cost by 60%, that is for $\delta = 1.6$, per user variable profits from transfers do not change, keeping in mind that variable profits from cash-in-cash-out transactions are constant because we did not alter $\mathbf{p}^{C,S}$. This price change increases the expected amount of transferred money by 9%, leading to an increase in consumer welfare.³¹ This amount is a conservative estimate of the possible

³¹In both models of peer-to-peer transfers and cash-in-cash-out transactions consumers have varying sensitivity, depending on their position on the demand curve. This introduces complications when pooling surpluses of individual customers when computing total welfare. Specifically, it is hard to compare utils of customers with low incomes with customers with large incomes. Consequently, any type of aggregation would imply particular weighting of welfare of these different groups. The usual solution is to rescale the individual surplus using dollar (or Shilling) units. This solution is appropriate for smaller industries with relatively

increase in welfare for the following reasons. First, we consider simple mark-up policies only; more sophisticated and complex policies would likely achieve better results. Second, the increase in expected consumer surplus would lead to more adoption, therefore the consumer surplus gain is likely to be higher in the long run. Third, more adoption should generate even higher total profits, which would enable the network to use a lower mark-up δ and maintain the same total profits.

Unfortunately, the above price discrimination policy is hard to implement, because setting the cash-out fee for transfers to zero and keeping the cash-out for transportation/storage unchanged is not incentive compatible. Specifically, in order to circumvent the cash-out fee, users may start executing P2P transfers even if their intent is to transport or store the money. Given the current technology of the network, we propose a limited and incentive compatible price discrimination schedule. In particular, the network could allow users to cash-out for free up to an amount received from a transfer, and charge a positive cash-out fee for all other cash-out transactions. To maintain incentive compatibility, the cash-out fee not preceded by a transfer has to be less than the transfer fee. Otherwise customers would have an incentive to transfer money to someone else in order to cash-out later with a lower charge. We find that if the network sets $\delta = 1.6$ and $\mathbf{p}^{C,S} = 0.75\mathbf{p}^T$ it obtains the same cash-in-cash-out profit as before, while increasing the transacted amount by approximately 10%. The above exercise shows that abolishing cash-out fees for users that conduct peer-to-peer transfers and having a cash-out fee for the rest of the transactions as described above is a Pareto improvement in comparison to present pricing.

Another reason why cash-out fees are set high may be to lock the consumers into the network. A successful lock-in could create the network effect which would lengthen the lifetime of the Shilling inside the network. The potential goal of the network is to charge fees to induce use patterns of m-money that are similar to a traditional monetary system in which money is reused extensively, as opposed to being a one time peer-to-peer transfer vehicle. In effect, such a change would lead to shifting the network's revenue to peer-to-peer homogeneous customer base, but in our case it drastically overvalues welfare of extremely affluent customers (with near zero price coefficients). We think that such re-weighting is not particularly desirable so we decided to use a total amount of transacted money as a measure of welfare.

transfers, which carry zero marginal cost to the company, as opposed to cash-outs. The welfare impact of such developments are ambiguous and are beyond the scope of this paper. However, the current pricing policy, despite large cash-out fees, has failed to induce more hops of a Shilling within the network; therefore, to induce more hops, a much larger cash-out fee may be required. Higher cash-out fees could lead to less usage overall, potentially decreasing profits. We conclude that it may be necessary to use different mechanisms to induce longer lifetime of the Shilling in the network, such as paying consumers interest on the balance,³² or increasing the density of agent network.

7.4 Willingness to pay for money transportation and storage

The transport-storage demand model allows us to quantify the consumers' willingness to pay to avoid carrying cash or storing cash at home. We construct a measure of the willingness to pay that is reminiscent of the compensating variation. In particular, to estimate the willingness the pay for avoiding carrying cash, we compare two transactions: (i) approximately 0 kilometers between cash-in and cash-out, within 1 day and at the current price schedule p ; and (ii) 'x' kilometers between cash-in and cash-out, within 1 day and at the counterfactual price schedule $p + y$. We find the willingness to pay y such that the expected amount of money transacted is the same in (i) and (ii).³³ The results are depicted in Figure 21. We find that consumers are willing to pay 400 Shillings extra not to walk an extra kilometer carrying cash. This amounts to roughly 1% of the average transfer size, which can be translated to a belief that walking an extra kilometer increases the probability of being robbed by 1%, if the users are risk neutral. This number is large, but it reflects the numbers cited in the survey by United Nations Office on Drugs and Crime (2009). This last calculation should be treated as an upper bound because it assumes that users are risk-neutral. If users are risk averse, the rationalizing belief of users about the probability of robbery would be smaller.

Similarly, to quantify the disutility of keeping savings at home we compare two transactions: (i) 0.5 kilometers between cash-in and cash-out, within approximately 0 days and at

³²In 2014, after the period of our data, Tigo introduced a dividend payment on the mobile accounts imitating positive interest rates. Tigo is the first mobile network to do so.

³³We choose transactions with one day time span to exclude savings motives from our estimates.

the current price schedule p , (ii) 0.5 kilometer between cash-in and cash-out, with a network lifetime of ‘ x ’ days and at the counterfactual price schedule $p + y$.³⁴ We find the willingness to pay y such that the expected amount of money transacted is the same in (i) and (ii). The results are presented at Table 22. We find that users are willing to pay 414 Shillings (about 1.1% of the average transfer) for an extra day of secure storage. This number should be interpreted as the difference between the perceived security of m-money and storing the money at home, and should be regarded as an upper bound.

The official criminal records in Tanzania, and similar African countries, are known to be extremely unreliable. For this reason, the usual way to measure the extent of crime are direct surveys, which are likely to be subject to a large survey bias. According to our knowledge, we provide the first estimates of criminal activity in Africa based on consumer behavior instead of using official records or surveys.

8 Conclusion

We analyzed a m-money network in Tanzania that is widely used to make (i) money transfers across users (P2P transfers); (ii) transport money for the same user without a transfer; and (iii) store money for the short- and medium-term as a savings account. We analyze and quantify the demand for P2P transfers in terms of distance, lifetime in the network, and number of transfers (hops) of a Shilling before cash-out. We find that the elasticity of demand for P2P transfers is decreasing in distance between cash-in and cash-out. We also find that the senders internalize to a significant extent the cash-out fees that the network imposes on receivers when money is cashed-out. Despite high cash-out fees, most users cash-out after only one transfer.

We find that a significant percentage of transactions on the network do not involve transfers, but are deposits and withdrawals by the same user despite the high cash-out fees. The network is used for storage, even though the interest rate is zero. However, quite surprisingly, the network is also used extensively for very short term (less than 2 hours) transportation of

³⁴We choose the baseline distance to be 0.5 kilometers, because we find that for longer lifetime transactions the distance between cash-in and cash-out is close to this number.

cash because of extremely high crime.

Using the demand estimates we provide measures of willingness to pay to avoid carrying cash in the pocket when traveling as well as keeping cash at home. We estimate that consumers are willing to pay up to 1% of the transaction amount to avoid carrying it in the form of cash for each extra kilometer. Similarly, they are willing to pay up to 1.1% to avoid keeping money at home for an extra day.

We find that present pricing is inefficient. We propose an incentive compatible and Pareto superior price discrimination scheme that would set the cash-out fee to zero for money received through a transfer while setting the cash-out fee a bit below the transfer fee for money that was deposited and withdrawn by the same person.

References

- AKER, J. C. (2010): “Information from Markets Near and Far: Mobile Phones and Agricultural Markets in Niger,” *American Economic Journal: Applied Economics*, 2, 46–59.
- ALAN, S., B. E. HONORÉ, L. HU, AND S. LETH-PETERSEN (2011): “Estimation of panel data regression models with two-sided censoring or truncation,” Tech. rep., Working Paper, Federal Reserve Bank of Chicago.
- BERRY, S., J. LEVINSOHN, AND A. PAKES (1995): “Automobile prices in market equilibrium,” *Econometrica: Journal of the Econometric Society*, 841–890.
- COLE, S., T. SAMPSON, AND B. ZIA (2011): “Prices or knowledge? What drives demand for financial services in emerging markets?” *The Journal of Finance*, 66, 1933–1967.
- DEMIRGÜÇ-KUNT, A. AND L. KLAPPER (2012): “Financial inclusion in Africa: an overview,” *World Bank, Washington, DC*.
- ECONOMIDES, N. (1991): “Compatibility and the creation of shared networks,” *Electronic Services Networks: A Business and Public Policy Challenge*, 39–55.
- (1996): “The economics of networks,” *International Journal of Industrial Organization*, 14, 673–699.

- FARRELL, J. AND G. SALONER (1992): “Converters, compatibility, and the control of interfaces,” *The journal of industrial economics*, 9–35.
- HONORÉ, B. E. (1992): “Trimmed LAD and Least Squares Estimation of Truncated and Censored Regression Models with Fixed Effects,” *Econometrica*, 60, 533–65.
- (2008): “On Marginal Effects in Semiparametric Censored Regression Models,” Princeton University.
- JACK, W. AND T. SURI (2013): “Risk Sharing and Transaction Costs: Evidence from Kenya’s Mobile Money Revolution,” *American Economic Review*.
- KATZ, M. L. AND C. SHAPIRO (1994): “Systems competition and network effects,” *The Journal of Economic Perspectives*, 8, 93–115.
- KING, R. G. AND R. LEVINE (1993): “Finance and growth: Schumpeter might be right,” *The quarterly journal of economics*, 108, 717–737.
- LEVINE, R. (1997): “Financial development and economic growth: views and agenda,” *Journal of economic literature*, 688–726.
- MALLAT, N. (2007): “Exploring consumer adoption of mobile payments – A qualitative study,” *The Journal of Strategic Information Systems*, 16, 413 – 432.
- MILLER, G. AND A. M. MOBARAK (2015): “Learning about new technologies through opinion leaders and social networks: experimental evidence on non-traditional stoves in rural Bangladesh,” *Marketing Science*.
- NARASIMHAN, L., K. SRINIVASAN, AND K. SUDHIR (2015): “Editorial: Marketing Science in Emerging Markets,” .
- PAKES, A. AND D. POLLARD (1989): “Simulation and the Asymptotics of Optimization Estimators,” *Econometrica*, 57, 1027–57.
- PLYLER, M., S. HAAS, AND G. NGARAJAN (2010): “Community-Level Economic Effects of M-PESA in Kenya: Initial Findings,” *College Park, MD: IRIS Center, Assessing the Impact of Innovation Grants in Financial Services Project*.

- QIAN, Y., Q. GONG, AND Y. CHEN (2014): “Untangling searchable and experiential quality responses to counterfeits,” *Marketing Science*.
- SCHUMPETER, J. A. (1912): *The theory of economic development*, Cambridge, MA: Harvard University Press.
- SUDHIR, K., J. PRIESTER, M. SHUM, D. ATKIN, A. FOSTER, G. IYER, G. JIN, D. KENISTON, S. KITAYAMA, M. MOBARAK, ET AL. (2015): “Research Opportunities in Emerging Markets: an Inter-disciplinary Perspective from Marketing, Economics, and Psychology,” *Customer Needs and Solutions*, 1–13.
- SUDHIR, K. AND D. TALUKDAR (2015): “The “Peter Pan Syndrome” in Emerging Markets: The Productivity-Transparency Trade-off in IT Adoption,” *Marketing Science*.
- THE FINANCIAL INCLUSION TRACKER SURVEYS PROJECT (2013): “Mobile Money in Tanzania: Use, Barriers and Opportunities,” .
- UNITED NATIONS OFFICE ON DRUGS AND CRIME (2009): “Victimization Survey in Tanzania,”
<http://www.unodc.org/unodc/en/data-and-analysis/Data-for-Africa.html>.
- ZHANG, K. (2015): “Breaking Free of a Stereotype: Should a Domestic Brand Pretend to Be a Foreign One?” *Marketing Science*.

A Figures and Tables

Transfer fee	Transfer bands		Vodacom
	Tigo		
	Before Jan 24th	After Jan 24th	
10	200-999	200-999	
20	1,000-1,999	1,000-1,999	
30	2,000-2,999	2,000-2,999	
40	3,000-3,999	3,000-3,999	
50	4,000-9,999	4,000-4,999	200-9,999
100		5,000-9,999	
200	10,000-49,999	10,000-19,999	10,000-49,999
250		20,000-49,999	
500	50,000-299,999	50,000-299,999	50,000-299,999
1,000	300,000-499,999	300,000-499,999	300,000-499,999
1,500	500,000-1,000,000	500,000-1,000,000	500,000-1,000,000

Table 1: Change in the non-linear fee schedule for transfers. All numbers are in Tanzanian Shillings.

Cash-out fee	Cash-out bands	
	Before Jan 24th	After Jan 24th
500	1,000-9,999	1,000-4,999
600		5,000-9,999
800	10,000-19,999	
1,000	20,000-49,999	10,000-24,999
1,250		25,000-49,999
1,400	50,000-99,999	
1,500		50,000-99,999
1,800	100,000-199,999	
2,000		100,000-199,999
3,000	200,000-299,999	200,000-299,999
4,000	300,000-399,999	300,000-399,999
5,000	400,000-1,000,000	400,000-1,000,000

Table 2: Change in the non-linear fee schedule for cash-out.

Unique customers	1400644
Mean transfer size	38751
Std. dev. of transfer size	83954
Fraction of day-senders with no transfer	95.5%
Fraction of day-senders with 1 transfer	3.7%
Fraction of day-senders with 2 transfers	1.1%
Fraction of day-senders with 3+ transfers	0.02%
Fraction of day-receivers with no transfer	95.0%
Fraction of day-receivers with 1 transfer	4.7%
Fraction of day-receivers with 2 transfers	0.39%
Fraction of day-receivers with 3+ transfers	0.01%
Mean transfer distance (km)	69.2
Std. dev. of transfer distance (km)	109.5

Table 3: Aggregate statistics about P2P transfers in the Tigo network.




Figure 1: Distribution of the amount transferred in Tanzanian Shillings.

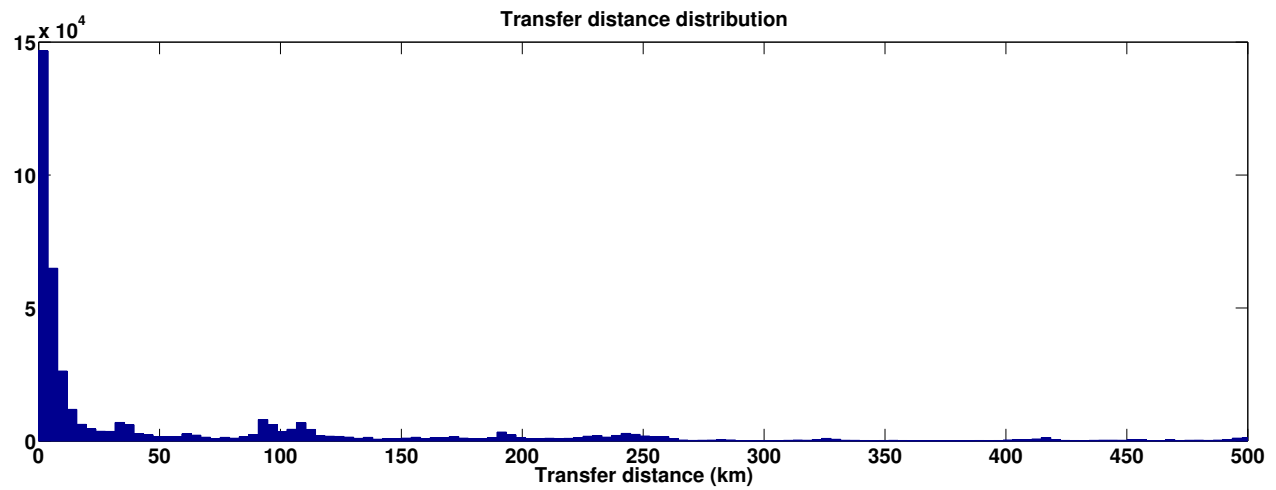


Figure 2: Distribution of the geographical distance of transfers.




Figure 3: Logarithmic heat map of number of transfers by origin.




Figure 4: Graph representation of first 1,000 transfers in the data.




Figure 5: Distribution of the lifetime of a Shilling in the Tigo network from cash-in to exit.

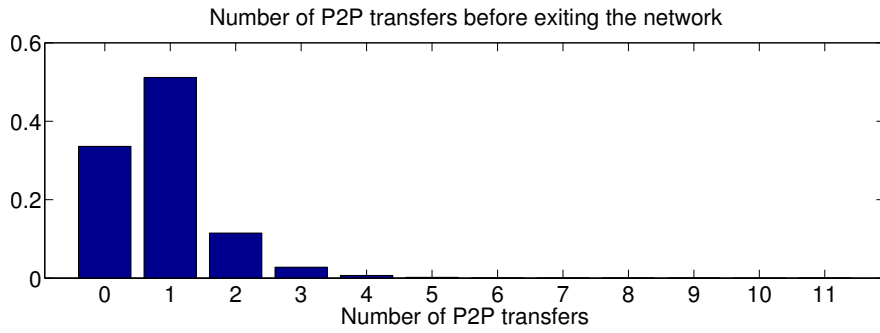


Figure 6: Distribution of number of P2P transfers during the lifetime of a Shilling in the Tigo network since cash-in.

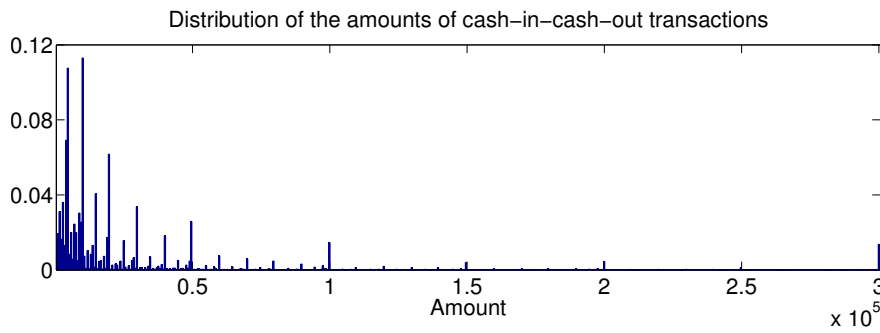


Figure 7: Empirical distribution of the cash-in-cash-out transaction sizes.




Figure 8: Distribution of the time between cash-in and cash-out for the transactions with zero hops.




Figure 9: Median distance between cash-in and cash-out for transactions with zero hops as a function of the lifetime that the money stayed in the network, for short lifetimes.




Figure 10: Median distance between cash-in and cash-out for transactions with zero hops as a function of the lifetime that the money stayed in the network, for long lifetimes.




Figure 11: Joint distribution of the time and distance between cash-in and cash-out for the transactions with zero hops.




Figure 12: Funding sources of cash-outs.

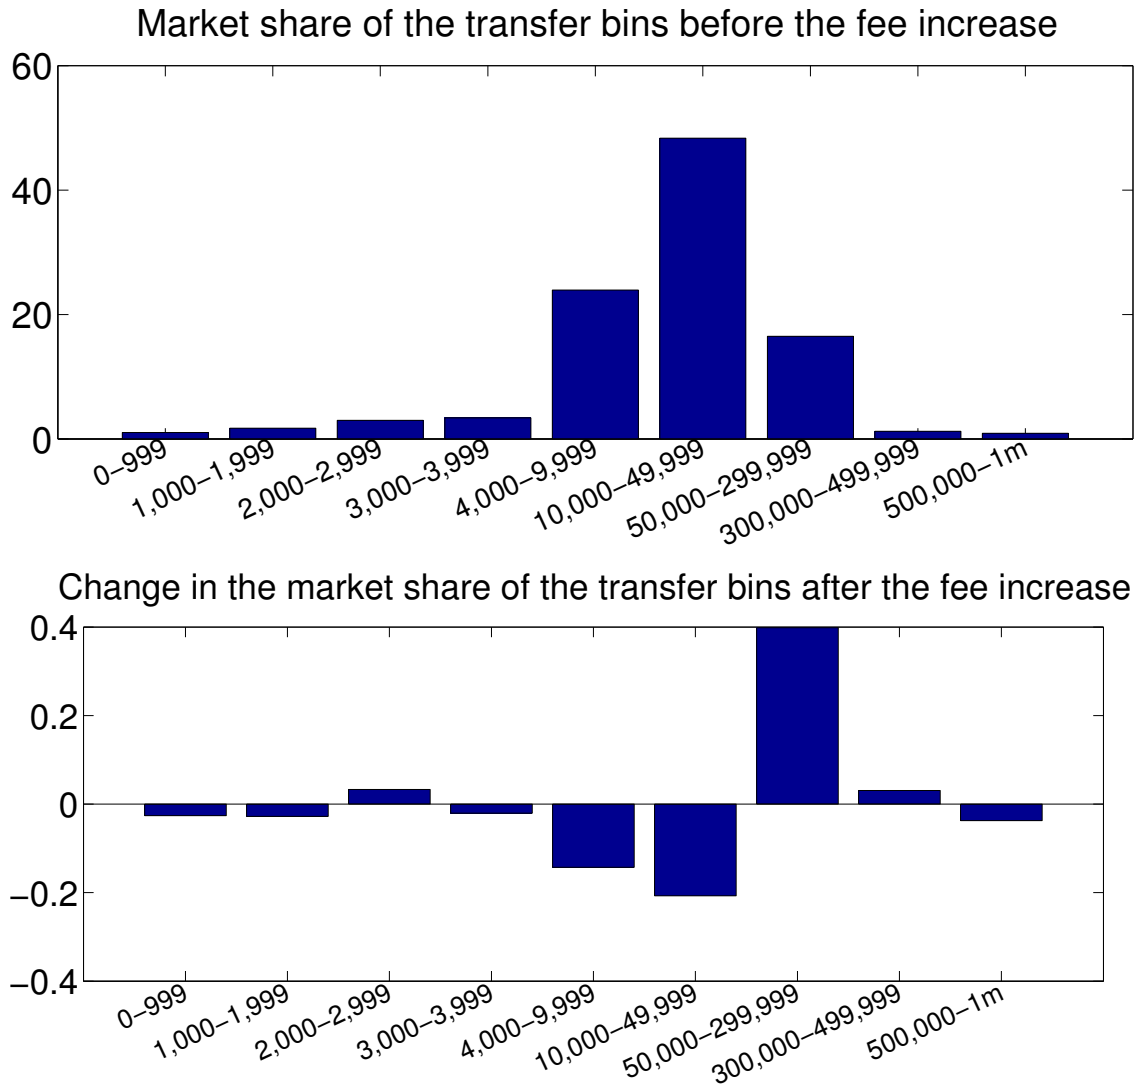


Figure 13: Market shares of transfer bins and impact of price increase.

Account balance (in thousands) -- before price change	8.64 (0.059)
Account balance (in thousands) squared -- before price change	-0.0024 (0.0001)
Account balance (in thousands) -- after price change	8.75 (0.049)
Account balance (in thousands) squared -- after price change	-0.0026 (0.0001)

Table 4: Effect of the recipient’s account balance on an average incurred daily cash-out fee. The estimation contains user fixed effects (not reported).

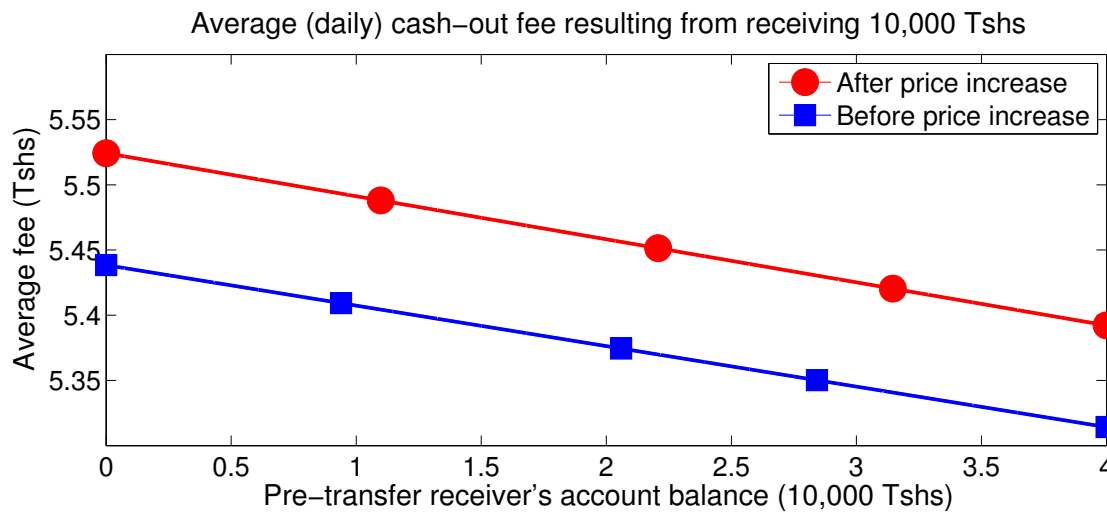


Figure 14: Increase in daily cash-out fees resulting from receiving 10,000 as a function of the pre-transfer balance of the receiver – before and after price change.

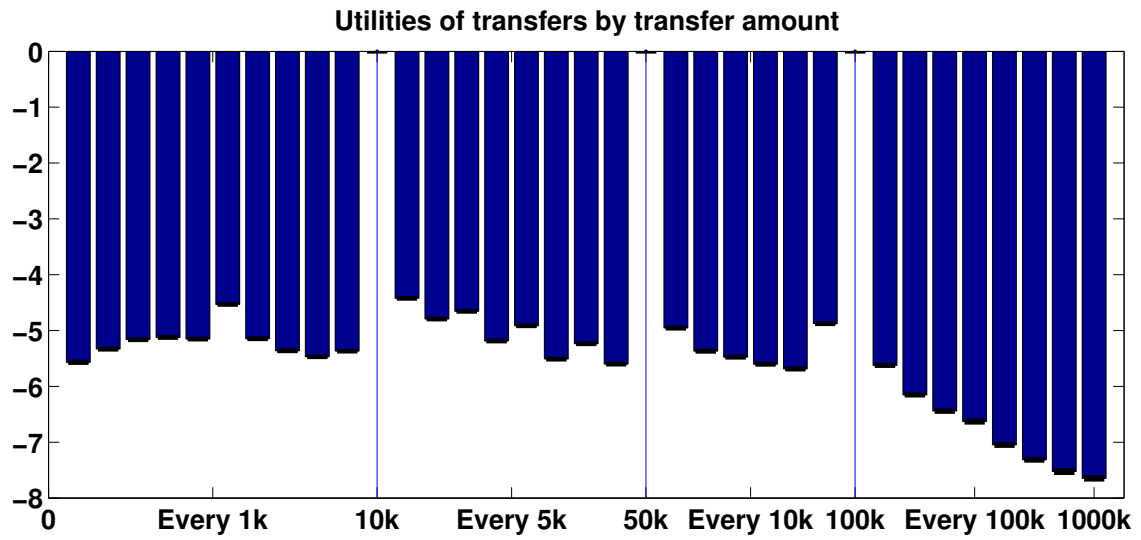


Figure 15: Fixed effects in the utility function grouped by the transfer amount.




Figure 16: Trend in the transfer utility.

Price coef. – mean ($\bar{\alpha}$)	-1.02 (0.018)
Cash-out coef. – mean ($\bar{\beta}$)	-43.62 (1.1896)
Price coef. – std. dev. (σ_α)	0.42 (0.014)
Cash-out coef. – std. dev. (σ_β)	21.2 (0.9)
Transfer size coef. – std. dev. (σ_κ)	0.2473 (0.0022)
Price coef./Cash-out coef. – Pearson correlation ($\rho_{\alpha\beta}$)	0.55 (0.046)
Price coef./Transfer size coef. – Pearson correlation ($\rho_{\alpha\kappa}$)	0.50 (0.026)
Cash-out coef./Transfer size coef. – Pearson correlation ($\rho_{\beta\kappa}$)	0.53 (0.017)
Covar. of ϵ for 2 adjacent transf. (ρ_ϵ)	0.07 (0.005)
Std. dev. of the persistent part of ϵ ($\sigma_{\bar{\epsilon}}$)	0.01 (0.008)
Std. dev. of the aggregate shock to the transfer util. (σ_A)	0.39 (0.006)
Mean of $\log(\text{balance}+1)$ before price change (μ_B^1)	6.35 (0.004)
Std. dev. of $\log(\text{balance}+1)$ before price change (σ_B^1)	2.33 (0.004)
Mean of $\log(\text{balance}+1)$ after price change (μ_B^2)	6.58 (0.004)
Std. dev. of $\log(\text{balance}+1)$ after price change (σ_B^2)	1.67 (0.005)
Distance coefficient (γ)	-0.28 (0.004)
Mean of $\log(\text{distance}+1)$ (μ_D)	1.08 (0.001)
Std. dev. of $\log(\text{distance}+1)$ (σ_D)	0.04 (0.007)

Table 5: Structural parameters of demand for peer-to-peer transfers.

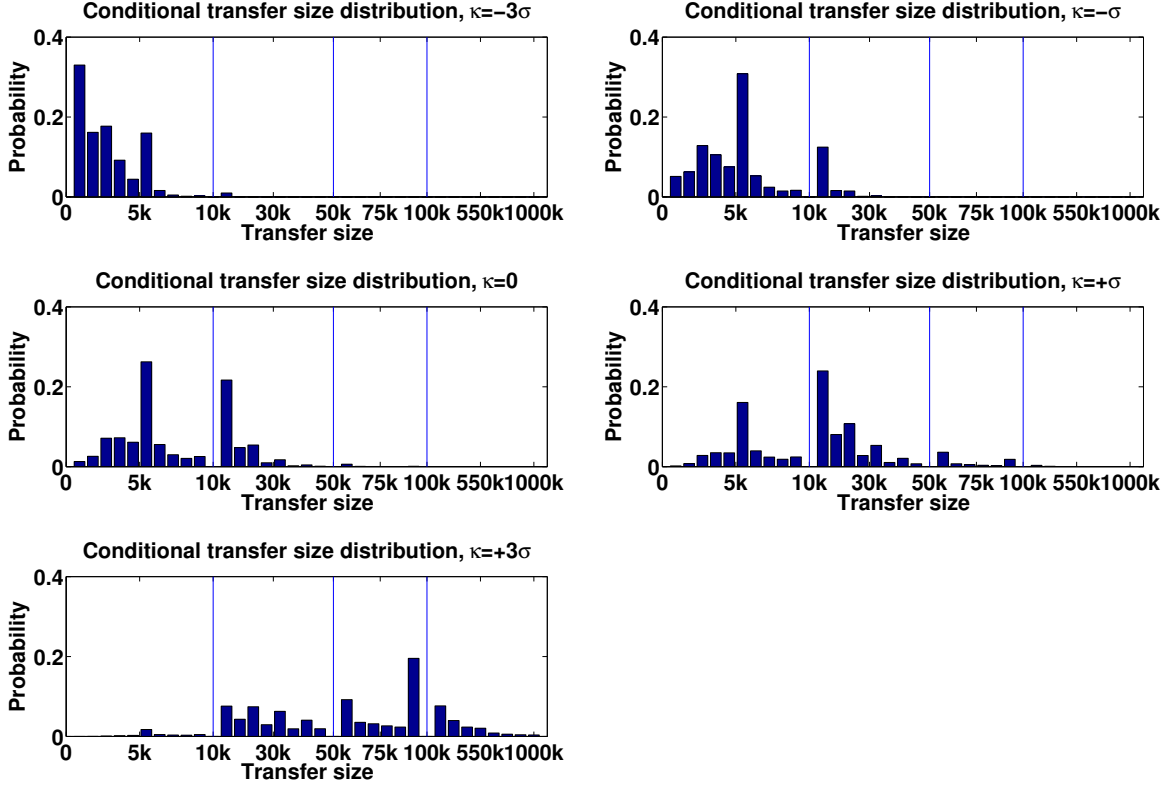


Figure 17: Transfer size distribution conditional on making a transfers for different values of κ .

	$\kappa = -3\sigma$	$\kappa = -3\sigma$	$\kappa = -3\sigma$	$\kappa = -3\sigma$	$\kappa = -3\sigma$
Transfer coefficient, α	-1.65	-1.22	-1.02	-0.81	-0.39
Cash-out coefficient, β	-77.27	-54.84	-43.61	-32.40	-9.96

Table 6: Mean of price coefficients conditional of different values of κ .

Peer-to-peer transfers	
Transfer price elasticity	-0.117
Transfer price elasticity – 0km transfers	-0.13
Transfer price elasticity – 5km transfers	-0.098
Transfer price elasticity – 10km transfers	-0.057
Transfer price elasticity – 15km transfers	-0.026
Transfer price elasticity – > 20km transfers	-0.01
Cash-out price elasticity	-0.35
<hr/>	
Cash-in-cash-out transactions	
Cash-out price elasticity	-0.88
Cash-out price elasticity – 0km, 0.5 day transactions	-0.82
Cash-out price elasticity – 5km, 0.5 day transactions	-0.80
Cash-out price elasticity – 0km, 5 day transactions	-0.81

Table 7: Price elasticities of demand.

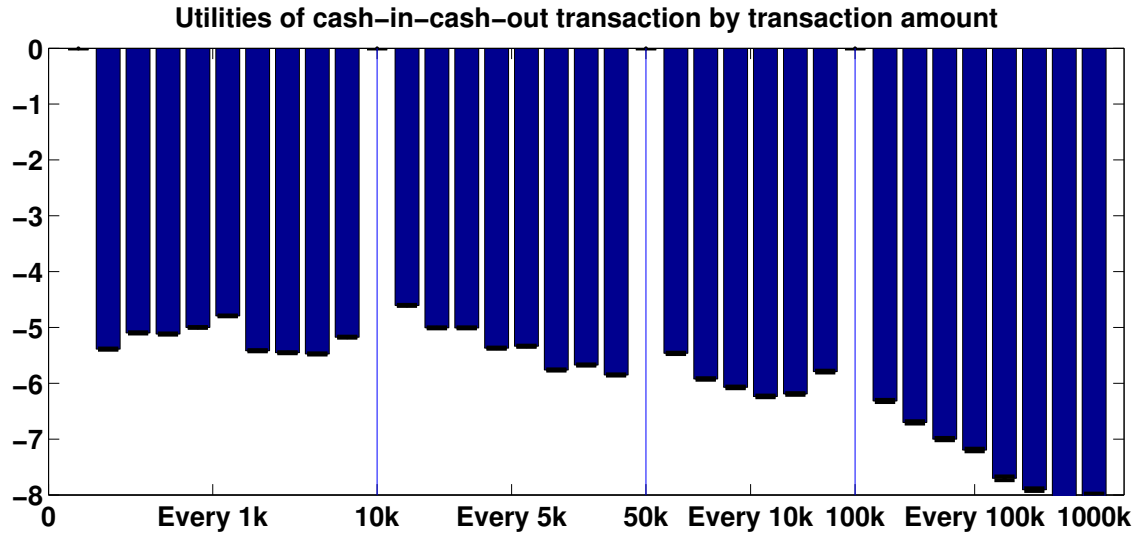


Figure 18: Fixed effects in the cash-in-cash-out utility function grouped by the transaction amount.




Figure 19: Trend in the cash-in-cash-out utility.

Cash-out price coef. – mean ($\bar{\beta}$)	-0.94 (0.011)
Cash-out price coef. – std. dev. (σ_{β})	0.36 (0.006)
Transaction size coef. – std. dev. (σ_{κ})	0.2172 (0.0018)
Price coef./Transaction size coef. – Pearson correlation ($\rho_{\beta\kappa}$)	-0.49 (0.016)
Covar. of ϵ for 2 adjacent trans. (ρ_{ϵ})	0.10 (0.002)
Std. dev. of the persistent part of ϵ ($\sigma_{\bar{\epsilon}}$)	0.01 (0.009)
Std. dev. of the aggregate shock to the transaction util. (σ_A)	0.39 (0.008)
Distance coefficient (γ)	-0.29 (0.004)
Mean of distance (μ_D)	11.69km (0.275)
Time-span coefficient (γ)	-0.28 (0.006)
Mean of time-span (μ_S)	2.18 days (0.056)
Distance/time-span – Pearson correlation (ρ_{DS})	0.00 (0.000)

Table 8: Structural parameters of demand for cash-in-cash-out transactions.

Effect on demand of both fees' increases	-3%
Effect on demand of transfer fee increase only	-2.5%

Table 9: Pricing counterfactual: impact on demand of the fees increases of January 24, 2013.

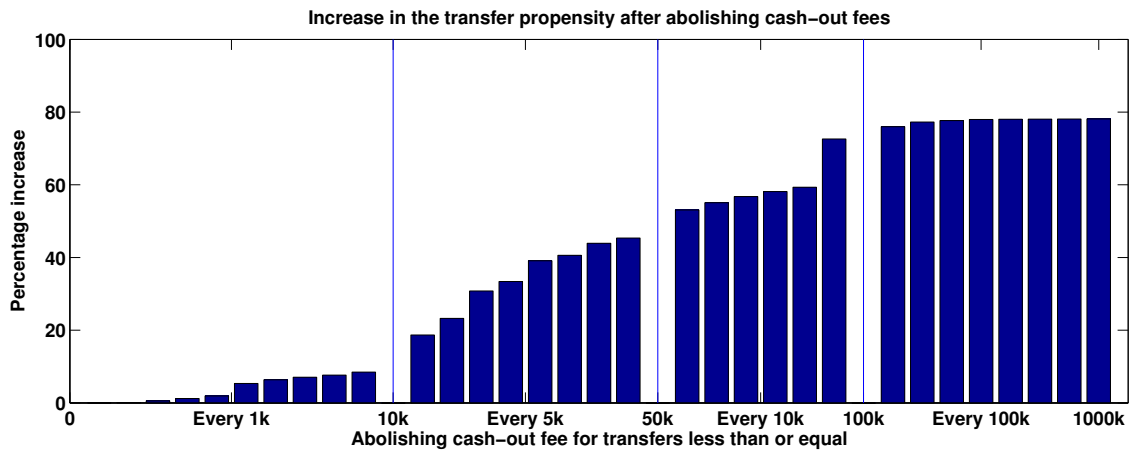


Figure 20: Percentage increase in the number of transfers after abolishing cash-out fee for users that make at least one transfer. The graph presents an increase for abolishing the cash-out fee for transfers smaller or equal than the value on an x-axis.

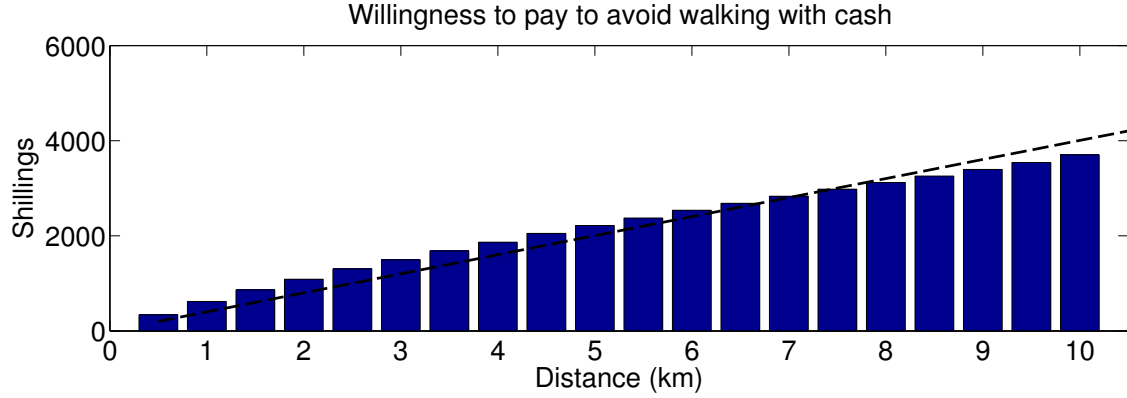


Figure 21: Willingness to pay to avoid walking with money. We compare two zero-hop transactions: (i) approximately 0 kilometers between cash-in and cash-out, within 1 day and at price schedule p , (ii) ‘ x ’ kilometers between cash-in and cash-out, within 1 day and at price schedule $p + y$. We find y such that the expected amount of money transacted is the same in (i) and (ii). Y-axis is y and X-axis is x . Dashed line is the best linear predictor.

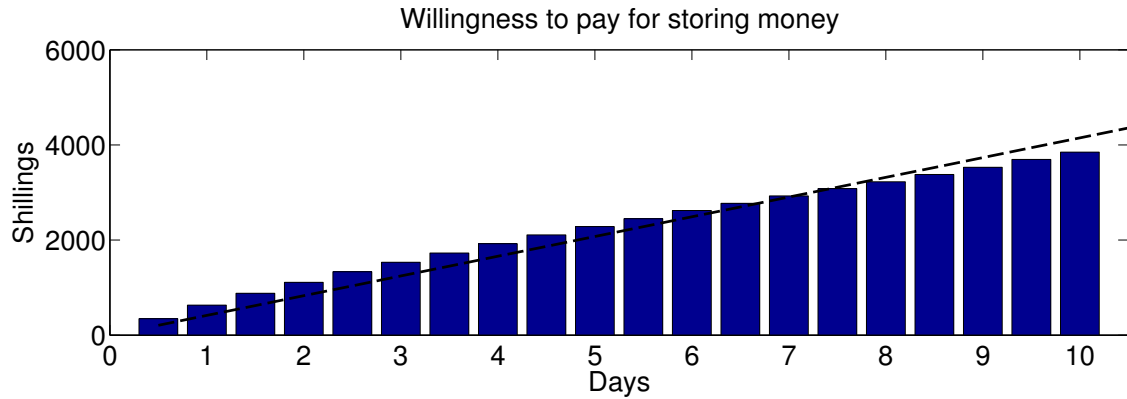


Figure 22: Willingness to pay to securely store cash over time. We compare two zero-hop transactions: (i) 0.5 kilometer between cash-in and cash-out, within approximately 0 days and at price schedule p , (ii) 0.5 kilometer between cash-in and cash-out, within ‘ x ’ days and at price schedule $p + y$. We find y such that the expected amount of money transacted is the same in (i) and (ii). Y-axis is y and X-axis is x . Dashed line is the best linear predictor.