

The Return to Capital for Small Retailers in Kenya:

Evidence from Inventories¹

Michael Kremer², Jonathan Robinson⁴, Olga Rostapshova⁵, and Jean Lee⁶

May 2010

VERY PRELIMINARY AND INCOMPLETE – PLEASE DO NOT CITE OR CIRCULATE
WITHOUT PERMISSION

Abstract

Standard textbook models suggest risk-adjusted rates of return should be equalized across activities within firms, and across firms. In general, measuring rates of return is difficult, but we take advantage of the characteristics of the retail industry to create estimates and bounds on the rate of return to inventories in a set of retail firms in rural Kenya. Using administrative data on whether firms purchased enough to take advantage of quantity discounts from wholesalers, we estimate a lower bound on rates of return for the median shop of 162 percent per year. A second approach measures the return to these investments by surveying shops on a regular basis about stockouts, or lost sales due to insufficient inventory. Though average returns are much lower from this approach (likely because firms fear losing customers due to lost goodwill), we are able to reject the hypothesis that the marginal rates of return are equal across shops. We examine alternative explanations of why we observe such high and heterogeneous marginal returns to capital and find evidence that credit constraints are not the only explanation, and our results suggest significant departures from the notion of firms as purely rational profit makers. (*JEL codes:* O10, O12, O16, O17. *Keywords:* returns to capital, microenterprise).

¹We thank David Autor, Austan Goolsbee, David McKenzie, Sendhil Mullainathan, Kevin M. Murphy, Ben Olken, John Sutton, seminar participants at the World Bank Microeconomics of Growth conference, the RAND Corporation, the NBER Africa program, the CEPR and lunch participants at UC Santa Cruz for their helpful comments. Dan Bjorkegren, Elliott Collins, Sefira Fialkoff, Eva Kaplan, Anthony Keats, Jamie McCasland, and Russell Weinstein provided excellent research assistance. We thank Edward Masinde, Isaac Ojino, Julius Okoth, David Oluoch, Dominic Ouma, and Bernard Yaite for collecting the data.

²Department of Economics, Harvard University, NBER, Brookings Institution, BREAD.

⁴Department of Economics, University of California – Santa Cruz.

⁵Harvard Kennedy School.

⁶Department of Economics, Harvard University.

1. Introduction

Standard textbook economic models suggest that the risk-adjusted rate of return should be equalized across activities within a firm. If capital markets function well, rates of return should also be equalized across firms, both within and even across countries. While it is clear that various frictions interfere with perfect equalization of rates of return across firms, it is not clear how big the departures from this benchmark are, and which departures are most important.

In addition, it is often difficult or impossible to directly measure rates of return to capital, particularly at the margin. In this paper, we take advantage of the structure of the retail industry among a subset of Kenya retailers to measure the rate of return to inventory capital. We are able to identify specific investment opportunities that are available to retail firms and directly compute the return that could be realized from these investments. The data imply very high marginal rates of return on average, and provide evidence for economically and statistically significant heterogeneity in marginal rates of return across shops.

In our first empirical strategy, we calculate the bounds on the rate of return to investment in inventory for approximately 900 shops in Western Kenya using a complete database of shop purchases from a major distributor of retail goods. We infer bounds on the rates of return to investments in inventory that could be achieved if shops shifted the timing of their purchases to take advantage of quantity discounts offered by the distributor. If shops have investment opportunities that exhibit diminishing returns at least locally, the average return on these incremental investments will be a lower bound on the marginal rate of return. Preliminary estimates using this approach suggest that the median firm has unexploited investment opportunities that would yield a real rate of return of at least 162 percent annual.

In our second empirical approach, we directly measure the expected rate of return to an incremental investment in inventory for small retail firms in Western Kenya. We collected detailed

panel data on inventory decisions, sales, and stockouts (lost sales in which a customer asks for a product that is out of stock and does not accept a replacement) for a sample of 142 small rural retail firms in 16 towns in Western Kenya. By measuring daily stockouts over a period of several months, we are able to measure the probability that an additional unit of inventory would have been sold in a given time period, had the shopowner bought it at the beginning of the period. In this way, we are able to estimate marginal rates of return to inventory investment by calculating the expected marginal benefit from holding an additional unit of inventory (the markup multiplied by the probability that the marginal unit would sell during the relevant time period), and comparing this to the marginal cost of obtaining an additional unit (the wholesale price multiplied by the cost of financing).

We focus our analysis on cell phone top-up cards, for several reasons. First, phone cards have fixed wholesale and retail prices and negligible storage and depreciation costs, and are not substitutable across brands. Second, phone cards are kept behind the counter in the shops we survey, so lost sales can be measured. However, one problem with this approach is that the costs of stocking out of phone cards might be quite high – shops may lose customer goodwill, and some evidence suggests that phone cards are “leader” items which get people into shops. Likely for this reason, we find modest returns to holding additional phone cards: the average shop in our sample would achieve a real rate of return of 39 percent to a marginal increase in inventory, with 8 percent of shops having returns greater than 50%. These returns are much higher than rates of return on debt or equity in either Kenya or international markets. If lost customer goodwill or other sales of complementary goods are significant, this will be a lower bound on the rate of return.

We explore the extent to which these rates of return may reflect high rates of return to capital or behavioral anomalies by separately estimating rates of return to two different brands of

phone cards, Celtel and Safaricom, for each shop in the sample. We present some preliminary tests of equalization of marginal rates of return across products within shops. On average, the rates of return for the Celtel and Safaricom products differ, although this appears to be driven by the top decile of the distribution of return. The median rates of return on these products are similar, and we find a rank correlation of 0.30 between rates of return for products of different brands.

If one treats the calculated rates of return as estimates rather than bounds, or assumes that all these bounds are equally tight because the cost of lost goodwill and other sales is similar across shops, we can then test whether these marginal rates of return are equal across shops, and estimate the degree of heterogeneity in rates of return under some assumptions about the underlying distribution of rates of return. Using a variety of tests, we reject the hypothesis of equalization of marginal rates of return across shops, suggesting some misallocation of capital in these markets. We find evidence that the standard deviation of the population distribution of annual rates of return may be as high as 98 percent.

This paper contributes a novel piece of evidence to a growing empirical literature on marginal rates of return to capital. Lucas (1990) famously noted that the simplest calibration exercise assuming a common aggregate production function suggests that the marginal rates of return to capital must differ dramatically between the rich and poor countries of the world. A recent paper by Caselli and Feyrer (2007) argues that the aggregate country level data on capital share of income, output, capital stock, are consistent with equalization of financial marginal rates of return across countries, after accounting for payments to previously unobserved factors (such as land and natural resources) and differences in prices of investment goods across countries.

The development literature, in contrast, finds evidence for high and variable marginal rates of return to capital. The approaches in this literature are varied and creative, but in general they find annualized marginal rates of return between 30 and 1200 percent, well above typical estimates for

the developed world. These studies fall roughly into three categories: revealed preference arguments, cross-sectional production function estimates, and evidence from exogenous shocks to credit access (in the form either of natural experiments due to policy changes or field experiments).

The first approach, as in Aleem (1990), notes that marginal rates of return must exceed the high interest rates at which people and businesses are willing to borrow, but may include some borrowing to smooth consumption as well as borrowing for productive investments .

The second method, which is employed in some form by much of the existing literature (as reviewed in Banerjee and Duflo (2004)), uses cross-sectional firm level accounting data to estimate production functions and infer rates of return from the estimated coefficients. These studies typically find evidence of high rates of return: Anagol and Udry (2006) find an annual rate of return of 150 to 250 percent to pineapple cultivation in Ghana, while McKenzie and Woodruff (2006) estimate that rates of return for entrepreneurs in Mexico with assets less than \$200 exceed 15 percent per month. However, while they provide an informative characterization of the economy, these cross-sectional estimates do not provide estimates of marginal rates of return.⁷

Finally, the third strategy exploits natural experiments or randomized field experiments to estimate marginal returns. Banerjee and Duflo (2005) examines policy shocks to directed lending in India and concludes that marginal rates of return to capital exceed 70 percent for those firms affected by the changes. Exploiting county-level variation in credit supply due to the Community Reinvestment Act, Zinman (2002) estimates gross rates of return to capital in the US on the order of 20-58 percent per year. Finally, de Mel, McKenzie and Woodruff (2009) estimate marginal rates of return of 60 percent for microenterprises in Sri Lanka in a field experiment in which the researchers

⁷The older estimates in this literature are often estimates of the average rate of return rather than marginals. All are subject to biases of an indeterminate sign and magnitude because they cannot separate returns to observed factors such as capital and labor from returns to unobserved factors, such as entrepreneurial ability. If higher ability entrepreneurs have been more successful and have accumulated more capital in the past, these rates of return will be biased upward. On the other hand, if credit market distortions prevent the efficient reallocation of capital to higher ability entrepreneurs (as in the knitted garment industry in Tirupur), higher capital firms may have lower ability and cross-sectional estimates of the rate of return to capital that ignore differences in ability may be biased downward.

provided grants or equipment valued at approximately one third of annual profits to randomly selected entrepreneurs.

In a recent study, Anagol and Udry (2006) take the elegant approach of using data on prices of used car parts of varying expected lifetimes to infer a 60 percent annual discount rate for taxi drivers in Ghana, although as they note, their estimate may not be directly interpretable as an estimate of the rate of return in a world with imperfect financial markets.

This paper is organized as follows. Section 2 discusses the context of the small-scale retail sector in Kenya. Section 3 describes the distributor data and the estimated rates of return implied by the bulk discount method. Section 4 introduces the stockout survey and data, presents the rates of return implied by the stockout data, suggests a framework for interpretation and shows that these data imply very high marginal rates of return for a nontrivial fraction of firms. Section 5 describes survey data on demographic, personality and entrepreneurial characteristics of the shop owners and describes their relationship to the estimated rates of return. Section 6 concludes.

2. The Small-Scale Retail Sector in Kenya

The small-scale retail sector comprises a significant share of economic activity in Kenya, particularly in rural areas. Daniels and Mead (1998) estimate that small and medium enterprises with 10 or fewer employees (not including agriculture and mineral extraction industries) comprise 12-14% of total Kenyan GDP, and that a quarter of this contribution comes from the retail trade.

We focus on a category of retail shop in Western Kenya called *dukas* in Kiswahili, which are small scale enterprises that are typically owner-operated, often by women and those with some secondary education. These shops are ubiquitous in market centers and small towns in the region, and are often located in clusters: adjacent to or in close proximity to several competing shops. They typically sell a relatively homogeneous set of household products such as perishable and non-

perishable foodstuffs, soaps, detergents, cooking fat, sodas, phone cards, and other household items. This is a poor area, and people consume a pretty small set of goods, so there is almost no variation across shops in the types of products they carry. The decision about how much inventory to buy is the most important optimization problem that most of these shopowners face. Most of these goods have low depreciation, and the storage costs are often not substantial.

Products are kept behind a counter (and often behind a set of metal bars) and all transactions and transfers of goods are mediated through the store operator. This means that we potentially have information on stockouts, although people may see certain goods out of stock and not ask for them. Phone cards, however, are kept below the counter so that customers are unable to know that they are out of stock without inquiring with the shopkeeper. Operators deal with a number of suppliers for the different goods they sell, but typically do business with a single supplier for each type of good.

Many goods are delivered on a regular schedule. Distributors are based in larger, semi-urban towns, and deliveries are made several times a week, depending on the product. For shops that are located in or near these larger towns, it is possible to restock from the distributor immediately if a stockout occurs or is soon to occur. For this study, however, we selected shops that are located too far from their distributors to make travel for restocking profitable. In addition, shops are able to restock certain products at any time by purchasing from a wholesaler that is located nearby. The disadvantage of restocking from these wholesalers is that they usually offer a smaller discount from the retail price than do distributors.

In this study, we focus our bulk discount analysis on non-perishable food items and household goods (e.g. vegetable cooking fat, soup mix, soap, and margarine) and our stockout analysis on top-up cards for cellular phone service. These products differ in their typical method of distribution. For the shops in our sample, non-perishable food items and household goods are

purchased either from the distributor or from a wholesaler, while phone cards are purchased exclusively from distributors. For our first empirical strategy, we obtained data on purchases from a major distributor that supplies over 100 household goods to retailers within the geographic area. This distributor makes deliveries to each shop once a week and offers the retailers bulk discounts based on total purchase amount at the time of delivery. The discounts range from .5% on purchases greater than 5,000 Ksh (approximately \$67) to 1.5% on purchases greater than 10,000 Ksh (about \$133), substantial relative to typical retail markup of about 10%. We have data on shop purchases from 2005 to 2009.

Our second empirical method uses stockouts of phone cards. Phone cards are a high volume commodity and are carried by many shops. The wholesale and retail prices are fixed and the markup is approximately 4-5%. There are two brands of top-up cards which are specific to the major cellphone carriers in the region: Celtel and Safaricom. Each brand has several denominations of cards. Celtel cards come in 40, 100, 200, 300, 600, and 1200 Kenyan shilling (Ksh) denominations.⁸ A small number of shops also have a technology which allows them to sell cards in arbitrary denominations. Safaricom cards come in 50, 100, 250, 500, and 1000 Ksh denominations. The brands are not substitutable for each other, though there is substitutability across denominations within a brand. Because most consumers buy the smallest available denomination, there is rarely substitution across denominations in the event in which a shopowner runs out of inventory for the desired denomination. One feature of the phone card distribution system that complicates our analysis is that goods must be purchased in discrete order sizes. For example, cards must often be purchased in packs of ten. For this reason, we calculate the expected profit from holding an additional order of ten cards rather than the return to one marginal card. Future work will explore how this discreteness may affect the analysis.

⁸ The exchange rate was approximately 70-75 Kenyan shillings per US \$1 during the study period.

3. Bulk Discount Analysis

3.1 Distributor Data

We analyze sales data from a major distributor of retail goods in Western Kenya. These data contain detailed records of purchases between January 1, 2005 and 2009 for purchases that are less than 100,000 Ksh in value, a rough cutoff which excludes very large wholesalers. We observe the name of the shop, date of the purchase, the quantity purchased of each product, the unit prices, the actual prices paid, the Value Added Tax paid, and any discounts received for each purchase. The shop identifiers also include some geographic information.

During this period, the distributor supplied 160 different household goods. While goods such as eggs, bread, milk and a set of other household goods are distributed separately and not observed in these data, the products in our data (including cooking fat, soap, detergent, soup mix, etc.) appear to comprise a significant share of inventory for small retail shops in the region.

We restrict our analysis to shops that purchase at least 5,000 Ksh worth of goods from the distributor in the first month in the data and that make purchases over a period lasting a minimum of 8 months. There are 585 shops in the data that satisfy this requirement, although we only have sufficient data for a subset of 434 of these to perform our rough first-pass calculation of the rate of return that could be achieved by taking advantage of quantity discounts.

The average shop satisfying these inclusion rules makes 40.7 purchases in the data, and the average length of time between the first and last purchase in the data is 571.4 days. The average shop in the sample invests 20,706 Ksh (\$276) per month in products sold by this distributor (although note that the distribution is skewed). Summary statistics for the data appear in Table 1.

Shops receive a 0.5 percent discount if their total bill including VAT exceeds 5,000 Ksh, a 1 percent discount if their bill exceeds 7,000 Ksh, and a 1.5 discount if their bill exceeds 10,000 Ksh.

Again, this discount is substantial relative to typical retail markup of 10 percent.

3.2 Estimates of Marginal Rates of Return from Bulk Discounts

We use the availability of bulk discounts to infer a lower bound on the average marginal rate of return. Figure 1a shows that shops do respond to the availability of bulk discounts by trying to make purchases that just exceed the discount thresholds -- there are bumps in the distribution at the cutoffs. However, a substantial fraction of purchases also fall in the intervals just below the discount thresholds, and shops frequently forgo the discounts they could achieve by buying a larger quantity of goods up front. On the other hand, shops sometimes do get discounts for total purchase amounts just below the threshold and we observe these discounts. Figure 1b presents the distribution of purchase sizes excluding data within 150 Ksh of threshold.

We can use this data to calculate a rough bound on the interest rates. If a shop does not buy enough to obtain a discount, this implies that the cost of financing the incremental purchase is greater than the benefit from the discount. We calculate the rate of return that each shop could have realized had it bought goods earlier in order to obtain the bulk discount, given that it would have been able reproduce the same sales pattern going forward as the sales pattern that is empirically realized. Our analysis does not let us estimate the interest rates, but only allows us to bound the marginal rate of return -- the advantage however, is that there is less of a concern about selection into the sample, since this is administrative data.

For example, suppose that a shop makes a 4,500 Ksh purchase each month. Given an interest rate of r over a period of a month, their cost of borrowing to get to 5,000 Ksh would be $500 * r$. The benefit would be a discount of $0.005 * 5000$. If they are not borrowing to get to the 5,000 Ksh threshold, this implies that $500 * r > 0.005 * 5000$, or $r > 0.05$. A 5 percent rate of return over one month would be equivalent to an annual rate of 82 percent.

Some shops have low turnover and buy very few goods from this distributor. For smaller shops it takes a long time to ever buy to the first threshold, so they are more likely to buy smaller quantities from wholesaler rather than distributor. Therefore, we cannot bound their rates of return very tightly.

We thus restrict our sample to shops that purchase at least 5,000 Ksh of goods in the first month they appear in the data and appear in the data for at least 8 months. These shops are generally larger than other shops, and have been in operation longer. To the extent that there are diminishing returns to scale, this sample will have lower underlying rates of return than the unrestricted sample, and the bounds we present should be regarded as lower bounds on the distribution of rates of return for the entire population of retailers. To the extent that larger shops are likely to have been in operation longer, this sample will likely exclude new shops that may not yet have learned to take advantage of the discount, for whom we might calculate a spuriously high bound for the rate of return due to a lack of information about the discounts.

We then search for the date on which they make a purchase that is closest to the next discount threshold. Using subsequent purchases, we then calculate the rate of return they could achieve by increasing the size of the purchase order to meet the next discount threshold. Using this method, we are able to bound rates of return for 434 of the 585 shops satisfying our inclusion criteria.

We find evidence that rates of return for a significant fraction of the shops we study can be bounded at extremely high levels and observe high heterogeneity. A large percentage of shops have huge returns. Under baseline perfect information assumption, median shop has rate of return of 162% annually. If we assume that it takes 50 to 100 percent longer to sell items that would have ordered, implied annual return is between 56% and 87% . Even under 100% adjustment, 20% of shops have annualized returns over 500% per year. Under that adjustment, over 10% of firms have

returns < 0 (adjusted for inflation) .

Figure 2 details the distribution of the lower bounds that we can place on the maximum rates of return over the sample period for each shop. For 68 percent of shops, we can bound their rates of return above 50 percent annual. For 54 percent of shops, we can bound annual rates of return above 100 percent. For 24 percent of shops, those that make purchases very close to the discount threshold in our data, we calculate bounds on annual rates of return above 1,000 percent. It's important to note that some firms, approximately 4.4% of the sample, never actually purchase less than the maximum, so we exclude them. The data presented is calculated from 523 firms.

There are several caveats to this analysis that should be noted. First, as can be seen from the figures, there is some bunching of purchases just below the discount threshold. In talking to the distributor, this appears to be due in part to the fact that discounts are given by truck drivers on their rounds, and at times in the past the distributor's system for checking that discounts were given correctly was not very effective. In particular, truck drivers seemed to have some discretion in giving discounts to shops that were just below the eligibility threshold. We are still in the process of sorting this issue out, so these results should not be seen as definitive.

Second, we calculate very high bounds on rates of return at some point in time for these shops, but average rates of return across time may be substantially lower. Third, in the current version of the analysis, we do not account for uncertainty over which products will be in demand. Shops may delay purchasing products until some of the uncertainty becomes resolved. In the medium run, we plan to account for this by analyzing the expected returns to very simple investment rules -- for example, the return to increasing the purchase order by equal amounts for the three highest volume products. To the extent that shops have more information, and could have chosen a higher return bundle of goods to buy, this will be a lower bound on the rate of return they could have achieved by increasing the order to the next discount threshold. For now, we note that the

finding of high rates of return is not likely to be sensitive to this modification in the calculation -- making the crude assumption that shops would have taken 50 percent longer to sell the *ex ante* optimal product mix than the one they actually do purchase, the median shop in our data would still have a rate of return to this type of investment bounded above 76 percent annual. Under the even more conservative assumption that it would take 100 percent longer to sell off the *ex ante* optimal bundle, the median shop would have a rate of return bounded above 49 percent annual.

Overall, these numbers are roughly consistent with Udry and Anagol's (2006) estimate of the rate of return for non-pineapple crops in rural Ghana, and well above other estimates of the annual rate of return to capital (de Mel, McKenzie and Woodruff, 2006; Banerjee and Duflo, 2004), although both the context and sample composition differ significantly from those studies.

4. Estimating Marginal Rates of Return from Stockouts

4.1 Survey Data

Our second approach uses stock-outs of phone cards to estimate the marginal rates of return. The dukas for our sample were recruited from a census of small retailers in 11 small towns in Western Kenya: Bumala, Funyula, Matayos, Mayoni, Nambale, Rang'ala, Segal, Sidindi, Shibale, Ugunja, and Ukwala. Shops were eligible to participate in the survey if they sold telephone cards, although a small number of businesses that sold these products but operate primarily as wholesalers were excluded from the sample. In addition, we excluded a small number of larger retail outlets (supermarkets") because they allow customers direct access to goods, so that the shopkeeper would have a difficult time observing and reporting the number of customers lost to stockouts. In total, 104 shops were eligible to participate in the survey in these 11 towns. Fifty-one shops initially refused to participate in the survey, and 8 withdrew from the survey (attributing their wish to discontinue the survey to its frequency, length and repetitiveness). After raising the compensation for participation, we recruited

a larger sample of shops of an additional 106 shops to participate in the survey from August to December 2007. The overall participation rate in the expanded sample is 74 percent. Our results can only be considered valid for the subset of shops that agreed to participate in the survey. However, in order to demonstrate that rates of return are not equalized it is sufficient to show that the rate of return to a particular investment in a well-defined subset of firms differs from that in another set.

In total, the analysis to date includes data on 142 shops which were surveyed twice weekly about a set of 33 products for a period ranging from three months to one year.¹⁰ The survey collected information about the number of items sold that day, the last time the shop had restocked each item, and the number of customers who had been lost to stockouts for each product.

As noted above, we define the event in which a customer comes to ask for a product that is out of stock and does not purchase a substitute to be a “stockout”. Daily data on stockouts for each item were constructed by asking shopkeepers to retrospectively report stockouts for each day since the previous survey. For some products, customers may substitute to another size or brand. To account for this, shopkeepers were asked whether the last customer on each day that requested a product that was out of stock substituted to another size or brand, or left. It was quite rare for customers to substitute to other brands or sizes -- substitutions were reported in fewer than 6 percent of cases. In these cases, we set the number of stockouts for that day to zero. This may bias the estimates towards zero, since customers who originally substitute from higher denomination cards to lower denomination cards may buy only one of the lower denomination cards in the event of a stockout, so that even cases in which customers substitute to other brands or denominations may result in lost sales. We plan to gather detailed information on the exact purchases made by customers who request products that are out of stock in a subset of future surveys in order to assess

¹⁰ Shops in the pilot sample of 15 shops were followed for one year. Based on the data from the pilot, it was determined that three months of data would be sufficient to estimate rates of return for each shop.

the extent to which this rough cut of the data accurately captures the revenues lost due to stockouts.

Data on wholesale and retail prices for all goods were collected from the suppliers. Retail prices deviate somewhat from the prices reported by retailers for some products, but there are likely to be very few deviations in retail prices for phone cards, since the cards are printed with their value. Informal interviews with shopkeepers also indicate that deviations from the retail price are rare.

Since shops are visited by distributors at regular intervals, the relevant horizon over which shop owners decided how much inventory to hold is the interval between distributor visits. We thus aggregate the daily data to shop-product-distributor visit interval observations in order to impute the marginal rate of return. In order to construct the number of stockouts over one of these intervals, we must have data on the stockouts and on the exact dates of distributor visits. As a consequence, we drop observations that belong to an interval in which we cannot construct a complete history of stockouts; we also drop observations that fall in intervals of indeterminate length because the date of a past or future distributor visit is missing.

Table 2 displays summary statistics for the sample. We observe each shop for a total of 237 days on average. Stockouts occur about 6 times per month for the average firm. Figure 3 shows the distribution of stockouts on a day for phone cards, conditional on having a positive number of stockouts that day.

4.2 Empirical Methods

These data allow us to directly compute the expected rate of return to buying one incremental unit of inventory.

The net rate of return to holding an additional unit of inventory over the time between distributor visits can be expressed as:

$$r = \frac{(P^R - P^W) \Pr[\omega > x] - ((1 - \delta) + c)}{P^W}$$

where r is the marginal rate of return on the inventory investment, P^W and P^R are the wholesale and retail prices, respectively, ω is the number of customers who want to buy the product, x is the level of inventory, δ is the rate of depreciation and c is the cost of storage. The return to holding an additional unit is just the markup multiplied by the probability the marginal good sells less depreciation δ and the cost of storage c divided by the wholesale price. This calculation implicitly assumes that the firm values unsold cards at the wholesale price at the end of the period, less depreciation and storage costs.

Typically, it is difficult to measure rates of return, since the expected rate of return to an inventory increase depends not only on expected extra sales, but also on product depreciation, storage costs, the risk of theft, and the cross-elasticity of demand with respect to other products. For these reasons, the ideal product to study would be one for which depreciation, storage, and expected theft costs are minimal, and one which is neither a substitute nor complement for other goods sold by shops. For these reasons, we focus on top-up cards for pay-as-you-go cellular phone service, which do not depreciate and take up little storage space. If there is no depreciation and if there are no storage costs, the expression for the return reduces to:

$$r = \frac{(P^R - P^W) \Pr[\omega > x]}{P^W}$$

These assumptions are approximately true for phone cards, which do not depreciate other than

through inflation¹¹ and are sufficiently small that the storage costs are negligible. Though theft is possible, no store in our survey reported any theft in the past year. Note that in these stores, phone cards and all other goods (with the possible exception of sodas) are typically kept behind the counter, so that customers do not have access to them unless they request the goods from the shopkeeper. Shops sell no substitutes for these goods other than top-up cards of other denominations, since the cards are specific to cell phone networks. They are unlikely to be strongly complementary to other household goods, but shops may incur some losses of sales of other products if they frequently stock out of phone cards due to a loss of reputation if customers prefer to buy all of their goods in one place. In this case, the estimates we present should be viewed as lower bounds on the actual rate of return.

Taking into account the minimum order size, the marginal rate of return to holding an additional pack of cards over the period of the investment (the interval between distributor visits) in this context is then given by:

$$r_i(D) = \frac{E[\min\{\max\{N_{ijt} - N_{ijt}^*, 0\}, N_{ijt}^{\min}\} | N_{ijt}^*, D] \cdot (P_j^R - P_j^W)}{P_j^W \cdot N_{ijt}^{\min}}$$

where $r_i(D)$ is the marginal rate of return to the investment over the interval of length D days for shop i ; P_j^W and P_j^R are the wholesale and retail prices of product j , respectively; N_{ijt}^* is the optimal (and actual) number of units of product j in stock at the beginning of the period; N_{ijt} is the number of customers who come to the store to buy the product (so that

¹¹All of the rates of return in this paper are presented in real terms, unless otherwise noted. According to the Central Bank of Kenya, inflation in Kenya was 9 percent in 2005/2006.

$\min\{\max\{N_{ijt} - N_{ijt}^*, 0\}, N^{\min}\}$ is the number of stockouts, capped at the minimum order size); and

N_{ij}^{\min} is the minimum number of units in a purchase from the distributor.

If the length of distributor visit intervals were constant across shops and across time, we could directly compute the expected marginal rate of return over those intervals from our data. In practice, the distributor visit intervals vary both within and across shops. For example, if a distributor visits a shop on Tuesdays and Fridays every week, the data will consist of intervals of three days and intervals of four days.

Note that $r(D) = \exp(\mathbf{r}D) - 1$, where \mathbf{r} is the daily interest rate. One option would be to substitute $\exp(\mathbf{r}_i D) - 1$ for $r_i(D)$ in (1) and treat it as a moment condition, and then use a generalized method of moments approach to obtain an estimate of the daily rate of return, \mathbf{r}_i .

Instead, we Taylor expand $r(D)$ around $\mathbf{r} = 0$ to obtain an estimating equation that is linear in \mathbf{r} :

$$\begin{aligned} r(D) &\approx (\exp(\mathbf{r}D) - 1)|_{\mathbf{r}=0} + (D \cdot \exp(\mathbf{r}D))|_{\mathbf{r}=0} + H.O.T. \\ &\approx \mathbf{r} \cdot D \end{aligned}$$

Substituting this into equation (5) for $r_i(D)$ and rearranging, we obtain the following estimating equation:

$$\min\{\max\{N_{ijt} - N_{ijt}^*, 0\}, N^{\min}\} = \mathbf{r}_i \cdot \left(\frac{D \cdot P_j^W \cdot N_{ijt}^{\min}}{(P_j^R - P_j^W)} \right) + \mathcal{E}_{ijt}$$

where \mathcal{E}_{ijt} is the error term.

We estimate daily marginal rates of return for each shop using OLS, Poisson, and negative binomial regressions. Our benchmark estimates are the OLS estimates, although Poisson and negative binomial estimates which take into account the count data nature of the outcome variable are also shown in Appendix Table 1. We then transform these to annual rates of return.

The OLS and Poisson specifications have an attractive robustness property. Viewed as quasi-maximum-likelihood (QML) estimators, both the OLS and Poisson estimators are consistent even if the distributional assumptions are wrong, as long as the model for the mean of the outcome is correct. If each shop faces a constant marginal rate of return over time, the model of the mean will be correct by construction.

The negative binomial regressions may potentially be preferred to the Poisson regressions because the Poisson regression restricts the mean and variance of the data generating process to be equal, a restriction which is clearly not satisfied in these data -- the sample variance of stockouts is an order of magnitude larger than the sample mean. However, the negative binomial regressions are not robust to misspecification of the distribution, and this is a case where making the econometric model more flexible hurts robustness.

To begin to interpret these estimates, we then estimate rates of return separately for each shop and phone card brand, since the standard theory would predict that rates of return should be equalized across products.

If cards are independent of other goods or if all shops face the same reputation costs from stockouts, then we can test for and estimate heterogeneity in marginal rates of return across shops. We do so by first using standard Wald tests of equality based on both the robust covariance matrix and the bootstrapped variance-covariance matrix. However, this test is not invariant to nonlinear transformations and since it relies on asymptotic approximations, may not be appropriate given the small number of shop-product-distributor visit intervals we observe for some shops in our data.

Thus we also construct a nonparametric permutation test, in the spirit of a Fisher test, to check whether the observed distribution of estimands is consistent with what we would expect to observe in a world of equalized marginal rates of return. If marginal rates of return are equal across shops, there are no unobserved components of the marginal cost of holding an extra unit and no unobserved marginal benefits that vary across shops, and there is no autocorrelation in shocks to demand for a shop, then we can view the distribution of stockouts for all shops as the empirical distribution of residual shocks to demand for all shops.

Under these assumptions, we can generate distributions of the variation in estimated interest rates that would be realized if shops in fact faced the same interest rate and thus the same distribution of residual shocks to demand. We do this by randomly assigning shop-product-distributor intervals to artificial shops, and generating simulated distributions of estimated interest rates. We then compare the actual distribution of estimated interest rates to the simulated distributions. We generate simulated distributions of the variance of the estimated interest rates, the 90-10 spread, and the Wald test statistic and compare the statistics for the actual distribution to the simulated distributions. We calculate the probability that the observed distribution of estimated coefficients would be generated at random under the null hypothesis of equal marginal rates of return by comparing the actual statistics (variance, 90-10 spread, and the Wald test statistic) to the empirical distributions of those generated by randomly permuting the shop assignment.

This procedure is robust to some types of correlation in shocks to demand over time. For example, if shocks to demand follow an AR(1) process and shops know this, then they will adjust their expectations accordingly. As a result, the *residual* shocks to demand for each shop will be uncorrelated over time.

Finally, we estimate the degree of underlying heterogeneity in rates of return in the population by using a random effects model. We estimate the following model:

$$\frac{\min\{\max\{N_{ijt} - N_{ijt}^*, 0\}, N^{\min}\}}{\left(\frac{D \cdot P_j^W \cdot N_{ijt}^{\min}}{(P_j^R - P_j^W)}\right)} = \bar{\mathbf{r}} + \mu_i + \varepsilon_{ijt}$$

where $\bar{\mathbf{r}}$ represents the average rate of return in the sample and $\mu_i \equiv (\mathbf{r}_i - \bar{\mathbf{r}})$. The object of interest is the standard deviation of μ_i .

34.3 Results

The results of our empirical strategy, the regression of E[min(stockouts, minimum order size)] on shop*(markup/minimum order size) is shown in Figure 4, which presents estimates of real rates of return and the standard errors. The preferred estimated annualized marginal rates of return fall between 0 and 750 percent in real terms (Figure 4). The average shop faces an annualized real marginal rate of return of 39 percent, though the median shop in the data faces an annualized real marginal rate of return around .26 percent. 18.3 percent of shops have a return greater than 50%. Some rates of return are negative, since an estimated nominal rate of return of zero would imply a negative real rate of return. Standard errors for the regressions are robust, which gives the appropriate QML standard errors, and are clustered at the shop level.

The rate of return bound calculated for the median shop using the stock-out method is below the estimate of the rate of return for the median shop in the bulk discount data. This could reflect differences in sample composition, since the distributor dataset contains a much larger sample, and includes all shops that buy from the distributor, including those in smaller and more rural market centers, while the stockout survey data was collected primarily in the larger rural market centers. However, while these shops may be selected, it's worth noting that even if the marginal rate

of return for all of the shops that chose not to participate were unexceptional, this would still suggest major constraints. In addition, the stockout estimates reflect average rates of return over time, while the bounds calculated using bulk discounts could reflect temporary shocks to rates of return.

One possibility is that these calculations overestimate the rate of return because customers are willing to intertemporally substitute and return on a later date to purchase a card if a shop runs out of stock. However, in the context we study, such behavior on the part of consumers is not likely to be empirically relevant because there are always a number of competitors nearby (within one hundred feet) who carry the same product, and market level stockouts are rare. We plan to gather data on this directly by surveying both shopkeepers and customers.

Another possibility is that we have not properly accounted for the possibility of theft in our calculations. First, note that stolen cards can be reported to the wholesaler and refunded in the case of theft, limiting the losses to the retailer. In addition, stolen cards are identifiable by serial number (reported on the receipt) and are inactivated and rendered worthless once reported stolen, reducing the value of these goods to a potential thief. Consistent with these institutional features, theft of phone cards appears to be extremely rare.¹²

While the probability of theft is observed to be low, it could be the case that it is increasing sufficiently sharply in inventory to explain the observed frequency of stockouts. Two features of our data suggest that a high marginal probability of theft is unlikely to explain stockouts. First, there is a large range of shop size within our sample: the largest shops in our sample carry a total value of inventory orders an order of magnitude larger than the inventory orders of the smallest shops. However, given that the probability of being robbed is very close to zero for all types of shops, even if the probability of being robbed is monotonically increasing in inventory, then the marginal

¹² One large wholesaler/distributor reported that only one theft of phone cards had been reported among several hundred customers in four months (April through July of 2007).

increase in the probability of being robbed with respect to an increase in inventory must be low on average across the observed range of inventory. On the other hand, if shops can make investments in preventing theft, what we observe is equilibrium theft probability as a function of size. A second line of argument relies on the intertemporal variation in stock within shops. There is substantial variation in the value of inventory held by a shop over time, and both the probability of theft at times of high and low inventory are close to zero. However, within shops, the investments made in theft prevention technologies (quality locks or security guards) do not appear to adjust with the relatively high-frequency changes in inventory; thus, the effect of the marginal increase in inventory on the probability of theft must be bounded by something very close to zero, and will not substantially affect our results.

A third possibility is that these stockouts reflect collusion on the part of shopkeepers to each hold low levels of inventories, since it is clearly socially optimal for there to be shop level stock outs but not market level stock outs. The information structure makes it difficult to believe that shops jointly decide how much inventory to hold, given that shopowners do not observe each other's restocking decisions and the stochastic nature of stockouts would make it difficult to verify deviations from any agreement. Direct inquiries confirm this intuition. In addition, the skewness of the within-market distribution of rates of return suggests that shopowners do not collude to hold lower levels of inventory than they would in a decentralized equilibrium – the simplest models of collusion would suggest that all shopowners would agree to reduce inventory and thus that stockouts should be relatively evenly distributed across shops within towns, but in fact the distribution is quite skewed, with some shops frequently experiencing stockouts and others only very rarely. We plan to further explore the degree to which collusion and market structure may influence this measure of rates of return by examining the correlation between the estimated rate of return and the competitiveness of the local market, as proxied for by the number of very local competitors, for

example.

These high rates of return do not appear to reflect failures to optimize driven by inattention or any other factor that would result in mean zero measurement error. However, they could be driven by behavioral anomalies that lead to difficulties in saving or systematic mistakes in setting inventory levels. Behavioral anomalies may be especially relevant in the context of developing countries due to structure of labor market.

In order to begin to explore whether these high rates of return reflect behavioral anomalies or genuinely high rates of return to capital, we separately estimate rates of return implied by stockouts of Celtel and Safaricom products for each shop and compare these returns. Because the pure credit constraint view implies equalization of returns across alternative investments within a firm, firms with extremely high returns on one brand are more likely to have extremely high returns on the other. For this analysis, we restrict attention to the shops that carry both Celtel and Safaricom products and calculate the interbrand correlations of returns. Table 3 presents the results, which show that the rates of return are related within shops the rank correlation between the rate of return on Celtel products and the rate of return on Safaricom products is 0.30. Those with very high Safaricom returns have very high Celtel returns, and so there's a strong correlation, but there is less correlation for lower returns. For example there's very little correlation for people in the more "normal" range of returns. So the relationship between the two types of returns is weak for shops with smaller returns, for instance for those with returns of less than 75% per year (column 2 of Table 3). This could be suggesting that there is a type of shop that has massive returns to both types of cards. This correlation is consistent with maximization, but may reflect similar mistakes in optimization for both brands of phone cards. Some limitations of this analysis include measurement error, and different losses of consumer goodwill from different types of stockouts. In future work we plan to examine this result more carefully, and formally test and reject equalization of returns

across investments in Celtel and Safaricom brand cards, and also bulk discounts.

We also look at time trends within our data to gather evidence to differentiate between the two competing hypotheses. To collect data on stockouts, we visited shops multiple times a week and measured lost sales. If shops were not optimizing, this data collection process could change behavior, by making the stock-out frequency more salient for the shop keeper. We thus look at how the rates of return varies with time we've been following shops. Since shops started data collection at different times, we can also control for overall time trends. Table 4 shows the results, calculated by sample. The old sample is less representative, while the new sample is larger, includes more types of shops. To get a shop-specific measure of learning, we regress the daily rate of return on the time trend since start of project (7/2005) and on time since each shop started stockout survey. The overall trend captures changes in the market. We find that returns decrease over the duration of our study. We observe significant downward time trends for all shops since the first visit at a shop and the start of the project, and see larger effects in old sample. This suggests that visits to administer the stockout survey indeed changed behavior, and that shopowners appear to have learned to stockout less frequently. The size of the learning effects can be observed intuitively, as follows: from column 1, reduce the annualized rate by about 1 point for every month in data. This is big at the mean but not very big for the 16% or so of shops at the right tail. There is also evidence for some time trends in the market, of roughly the same order of magnitude as the learning trend. The presence of the significant downward time trends for all shops since the first visit at shop and the start of the project provides evidence for the hypothesis that shops are failing to optimize.

In future work, we plan to run additional tests of whether rates of return reflect optimization by using a difference-in-differences strategy to look at how stockouts respond to wholesale price changes, and by testing whether apparent discrepancies in rates of return across brands are larger for those who might be expected a priori to make more mistakes, such as shopkeepers with less

experience or less education.

We find evidence that while the marginal rates of return to these inventory investments is similar to interest rates in the population, we do find significant evidence that they are heterogeneous across shops. With a standard Wald test based on the robust covariance matrix or on the bootstrapped covariance matrix, we can reject the hypothesis that the estimated interest rates are equal across shops at the 1 percent level. Since the Wald test is not invariant to nonlinear transformations, we perform this test on the estimated coefficients, the daily rate of return, and the annual rate of return, and find similar results.

In addition, using the permutation test described above, we find that the standard deviation of the observed distribution of estimated coefficients falls in the 99th percentile of the simulated distribution of variances (Figure 3). At 123 percent, the actual standard deviation of the estimated rates of return falls far above what would be expected if the shops actually faced the same rate of return. Taken together, we interpret these tests as a strong rejection of the hypothesis that the marginal rates of return to these investments are equal for the shops in our sample.

Given that we can reject homogeneity of returns across shops, we next estimate the extent of the heterogeneity with a random effects OLS regression. Some of the variation in the distribution of fixed effects reflects sampling error, so we use a random effects model to estimate how much of this variation reflects real underlying heterogeneity in rates of return in the population of shops. We estimate that the standard deviation of the annual rate of return in the population is 98 percent.

The parametric assumption of normality in both the distribution of rates of return and the error term is almost certainly wrong, and the estimate of the extent of underlying variation changes dramatically with different assumptions about the distribution of random effects or with a random effects Poisson model. However, the conclusion that there is a large amount of underlying heterogeneity in the rates of return is qualitatively robust to the choice of specification -- the

estimates of the underlying real variation in rates of return are consistently large. This provides evidence for economically significant departures from the equalization of rates of return across firms that would be predicted by the standard model.

As noted above, these results on heterogeneity should be interpreted carefully, as there may be unobserved heterogeneity in the costs of stockouts (such as lost sales of other goods or reputation costs) that could explain some fraction of the differences in rates of return across shops. We plan to test for reputation costs by examining the cross-sectional relationship between the density of competitors in the immediate vicinity of the shop and the imputed rate of return, and also by using a difference-in-differences strategy to estimate the impact of entry and exit of nearby competitors on stockouts. While not interpretable as causal estimates, these correlations would provide some idea of whether reputation costs are likely to be empirically significant in this context.

One additional issue in the current set of results is that our sample included only roughly 75% of the shops operating in the towns that we studied. However, our key results on both the level and variance are qualitatively robust to sample selection issues. In addition, rejecting the hypothesis of equal rates of return in the sample we do observe is sufficient to reject the hypothesis of equal rates of return in a larger sample.

5. Background survey

We also collected extensive data on the characteristics of shops and their owners, for the shops in both the bulk discount and the stock-out samples. The shops were administered a detailed background survey which gathered information on a number of standard demographic and background measures such as the owner's age, sex, ethnicity, educational attainment, literacy, and the size of the owner's family. In addition, we collected exhaustive inventory data, documenting each product and its value present in the shop at the time of the survey. The survey also included

questions on the shopowner's access to savings and credit, land, durable good and other asset holdings, transfers he had given and received, other sources of income, self-reported credit constraints, and access to finance.. Since trade credit provided by suppliers may also potentially be an important source of financing, a separate section of this questionnaire focused on the relationship between suppliers and the retailer, especially regarding any credit provided by suppliers.

In addition, we developed survey instruments to measure less commonly studied correlates of business success, including attitudes towards entrepreneurship, various psychological measures, and a host of measures of cognitive ability. This data provide valuable information on the entrepreneurial environment, how entrepreneurs think about investment decisions, and which characteristics best predict business success. We have almost finished collecting this data for a sample of approximately 900 shops.

In order to estimate the relationship between owner characteristics and the annual rate of return, we ran variance weighted least squares correlations, presented in Table 5. We did the analysis using annual rates of return and log of rate of return. We found that most of the background characteristics did not correlate with the rates of return. One of the variables that did appear to be significantly negatively correlated with rate of return was owners stating they don't get the bulk discount because they can make more investing elsewhere.

6. Conclusion

We use evidence from inventories to provide a novel look at the marginal rates of return to investments available to rural retail enterprises in developing countries. With administrative data on a large sample of shops, we use bulk discounts to estimate a lower bound for the marginal annualized rate of return to capital for the median shop of 162 percent. Even adjusting these numbers to account for uncertainty in future demand, we still estimate a lower bound on returns of

56 to 87%, and find that 20% of shops have annualized returns of over 500% per year. Using our second (stock-out) method of estimation, we calculate the average marginal rate of return is 39% per year, with 8% of shops having returns greater than 50%. We also find evidence for substantial heterogeneity in marginal rates of return among these shops -- using several tests, we reject the hypothesis that the estimated marginal rates of return are equal across shops.

This suggests the potential gains from improving the allocation of capital may be large. The ability to realize these gains and the policy levers most conducive to doing so depend on the sources of these differences. There are of course multiple potential hypotheses about why we observe such high and heterogeneous marginal returns to capital and why rates of return will not be equalized. The first explanation is the hypothesis that credit constraints prevent small shopkeepers from borrowing to equalize returns with the outside credit market. The second is the hypothesis that behavioral factors limit the ability of small entrepreneurs to equalize rates of return across different items within their firms. Our preliminary analysis suggests that credit constraints are not the only explanation: we find evidence for learning over time and do not find shop or owner characteristics to predict returns. We also do not find any evidence that the returns in the stockout analysis are correlated with self-reported credit constraints or with assets, suggesting that access to finance is not the explanation for our results. Our findings strongly suggest that shop owners are not fully optimizing – the process of being surveyed seems to have changed their behavior since to measure stockouts, we visited shops regularly to ask about the number of lost sales they had. Shops experienced fewer stockouts the longer they had been visited (even controlling for seasonal fluctuations in demand). In ongoing work we hope to be able to provide more information to help differentiate between these hypotheses by looking at rates of return on different items, comparing rates of return across shops from the “phone card” test and bounds on rates of return from the reordering” test, and examining correlations between rates of return on inventories as we measure

them and other characteristics.

Evidence that some shops don't optimize is related to broader issues of development. In developed economies, firms make mistakes too, but efficient equilibrium is arguably quickly achieved as efficient producers expand while inefficient producers get replaced. This mechanism may not work as well in developing contexts, such as Kenya, due to labor monitoring problems; government expropriation risk; and retail price maintenance, including manufacturer facilitated price collusion to keep markup high, induce retailers to carry brand, where free entry can lead to excess entry. We plan to examine these issues in our future work.

We measure the rate of return to investment in a narrow category of activities, and this is sufficient to reject the standard model. Under stronger assumptions, the rate of return we measure also provides information about the rate of return to a broader set of investments. The marginal rate of return we measure may also reflect the marginal rate of return to capital in a broad swath of rural economic activities if the individuals we study (or the households to which they belong) are diversified and allocate their working capital across a set of productive activities (such as farming, raising poultry, etc). Diversification has important implications for the interpretation of the estimand not only for this reason, but also because if these shopowners are diversified, it may be possible to interpret this rate of return as the social marginal rate of return rather than just the private rate of return. Aggregate stockouts in these market towns are rare and there are typically many shops selling the same goods, so in the context of rural retail shops, the social return to financing the purchase of an additional unit of inventory by any one shop may be close to zero -- if a customer finds that one shop has stocked out of a particular product, he will buy from a competitor. However, if shop owners are diversified and participate in a variety of productive activities including some that do not exhibit this zero-sum feature, then the marginal rate of return we measure may reflect the social marginal rate of return to capital as well as the private return.

6. Bibliography

- Aleem, I. (1990): "Imperfect Information, Screening and the Costs of Informal Lending: A Study of a Rural Credit Market in Pakistan," *World Bank Economic Review*, 3, 329–349.
- Anagol, S., and C. Udry (2006): "The Return to Capital in Ghana," *American Economic Review Papers and Proceedings*, 96(2), 388.
- Banerjee, A., and E. Duflo (2004): "Do Firms Want to Borrow More? Testing Credit Constraints Using a Directed Lending Program," *Mimeo*, MIT.
- Banerjee, A., and K. Munshi (2004): "How Efficiently is Capital Allocated: Evidence from the Knitted Garment Industry in Tirupur," *Review of Economic Studies*, 71(1), 19.
- Caselli, F., and J. Feyrer (2006): "The Marginal Product of Capital," *Forthcoming*, *Quarterly Journal of Economics*.
- Daniels, L., and D. Mead (1998): "The Contribution of Small Enterprises to Household and National Income in Kenya," *Economic Development and Cultural Change*, 47(1), 45.
- de Mel, Suresh, David McKenzie and Christopher Woodruff (2008), Returns to Capital in Microenterprises: Evidence from a Field Experiment, *Quarterly Journal of Economics* 123 (4): 1329-1372.
- Lucas, R. E. (1990): "Why Doesn't Capital Flow from Rich to Poor Countries?," *American Economic Review Papers and Proceedings*, 80(2), 92.
- McKenzie, D., and C. Woodruff (2006): "Do Entry Costs Provide an Empirical Basis for Poverty Traps? Evidence from Mexican Microenterprises," *Economic Development and Cultural Change*, *Forthcoming*.
- Zinman, J. (2002): "Do Credit Market Interventions Work? Evidence from the Community Reinvestment Act," *Mimeo*, Federal Reserve Bank of New York.

Table 1: Summary Statistics, Distributor Data

	Mean	Variance	N
Average number of purchases	40.67	26.68	585
Average number of days between first and last purchase in data	571.35	155.17	585
Average purchases per month, Ksh	20706.21	56053.25	585

Table 2. Summary Statistics from Stockout Sample

Summary Statistics: Stockout Sample	
Years Education	10.92 (3.28)
Age of Owner	34.74 (9.58)
Owner is Male	0.48 (0.50)
Years Shop Open	6.92 (6.82)
Other income last week (USD)	18.45 (80.07)
Value of durable goods owned (USD)	2852.61 (6372.47)
Value of animals owned (USD)	804.95 (1074.81)
Got a loan from bank in past year	0.10 (0.30)
Got a loan from MF institution in past year	0.21 (0.41)
Has bank account	0.69 (0.46)
Can read English	0.95 (0.23)
Participates in ROSCA	0.48 (0.50)
ROSCA contributions in past year (USD)	101.82 (167.17)
Owns premises	0.80 (0.40)

Notes: data available for 93/117 shops in sample.
Standard errors in parenthesis.

Summary Statistics from Stockout Survey	
Total days covered	237.52 (178.00)
Distributor intervals covered	64.08 (50.19)
Average length of distributor interval	3.71 (1.68)
Number of Products carried	3.41 (1.60)
Average stockouts per month	6.65 (9.77)
Probability of Restocking per period	0.80 (0.40)
Notes: data available for 117 shops, and 7560 shop distributor intervals.	




Figure 1a: Distribution of Purchase Sizes in Distributor Data for Shops Satisfying Inclusion Criteria (buying at least 5000 Ksh a month), January 2005 to December 2006 (Ksh)






Figure 1b: Distribution of Purchase Sizes in Distributor Data for Shops Satisfying Inclusion Criteria, but excluding data within 150 Ksh of threshold, January 2005 to December 2006 (Ksh)



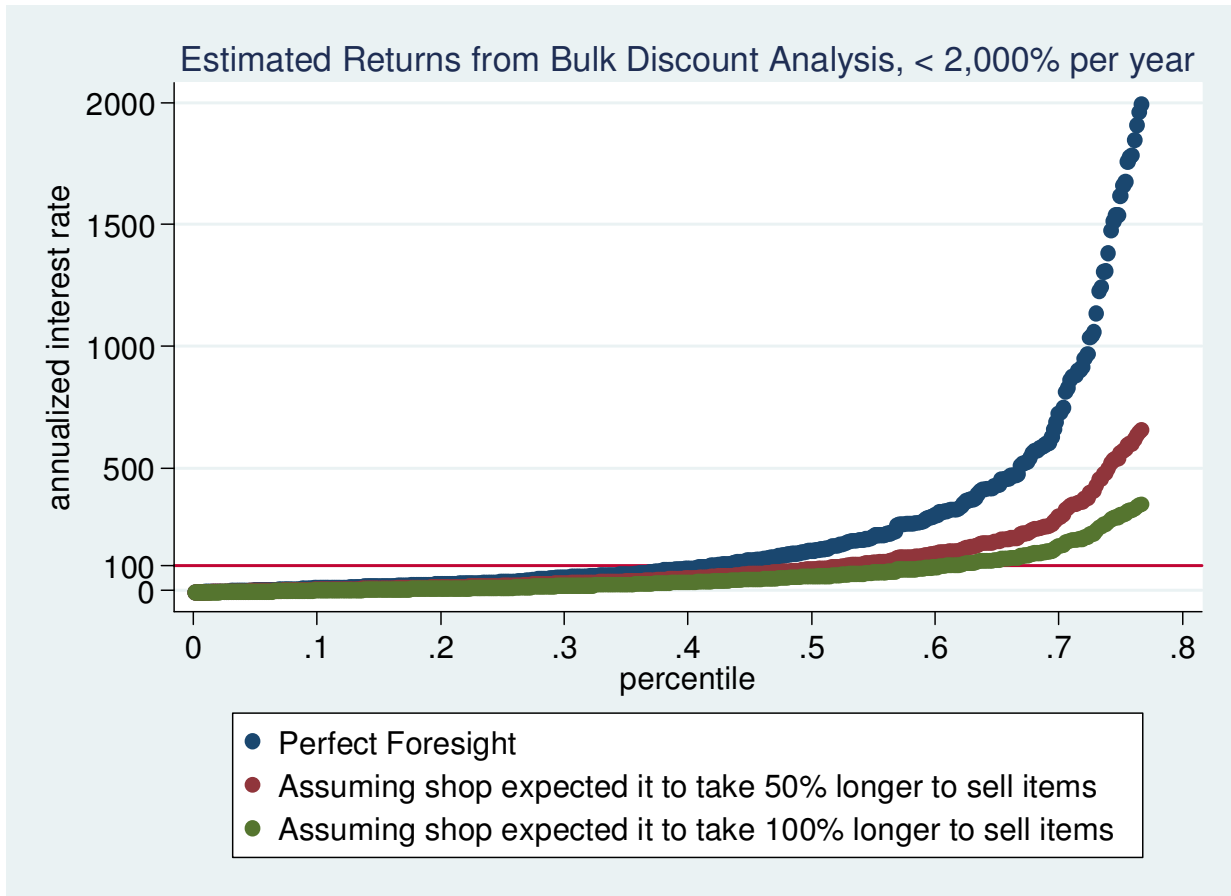


Figure 2: Distribution of Lower Bounds in Distributor Data. Panel A shows the distribution truncated at the 75th percentile; Panel B shows the distribution truncated at the 50th percentile; Panel C shows distribution truncated at the 80th percentile.




Figure 3: Distribution of Stockouts of Phone Cards Per Day, Conditional on Having a Positive Number of Stockouts


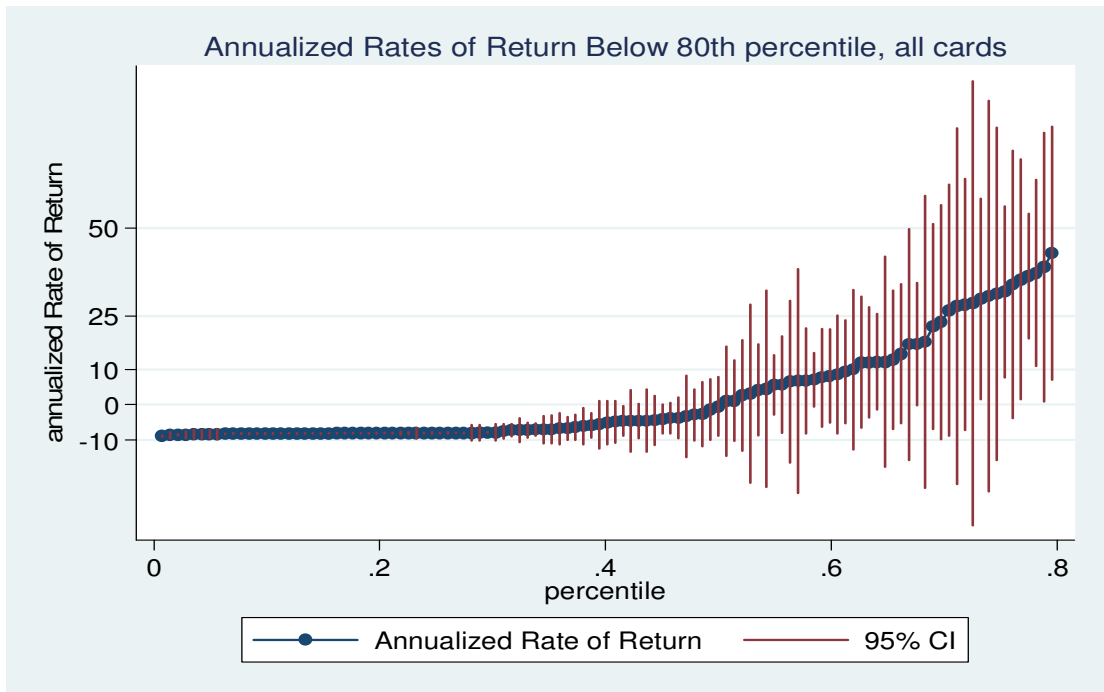


Figure 4: Distribution of Estimated Interest Rates from Stockout Survey Data (OLS Regression, Confidence Intervals Shown). Panel A shows the full distribution; Panel B truncates the distribution at the 80th percentile.




Figure 5: Distribution of Standard Deviation of Estimated Rates of Return for Simulated Shops Under Null Hypothesis of a Common Marginal Rate of Return, Stockout Survey Data. The line at 123 percent notes the actual standard deviation of estimated rates of return.

Table 3: Interbrand correlation of annualized return of Celtel and Safaricom phone cards

	(1)	(2)	(3)
	Annualized Return on Safaricom	Annualized Return on Safaricom	Dummy for Return on Safari > 50
Annualized Return on Celtel	0.297 (0.022) ^{***}	0.021 (0.027)	
Dummy for Return on Celtel > 50			0.297 (0.083) ^{***}
Controls for Card Denomination/Brand	N	N	N
Observations	89	78	89
R-squared	0.68	0.01	0.13
Spearman Correlation	0.55	0.46	0.00
Correlation	0.82	0.09	0.00
Mean of Independent Variable	415.803	4.988	0.171
Mean of Dependent Variable	68.35	21.68	0.21
Sample	All	Only those with returns < 75% per year	All

Notes: Column 1 excludes 1 individual with rate of return > 50,000 per year.

Table 4: Regression of daily interest rate on these independent variables

	(1)	(2)	(3)	(4)	(5)
Days Since first Visit at Shop	-0.050	-0.052	-0.040	-0.075	-0.075
	(0.006)***	(0.006)***	(0.004)***	(0.022)***	(0.022)***
Days since Project Start (July 2005)	-0.055	-0.06	-0.051	-0.187	-0.187
	(0.006)***	(0.007)***	(0.013)***	(0.027)***	(0.027)***
Days Since first Visit at Shop * New Sample					0.035
					(0.023)
Days since Project Start (July 2005) * New Sample					0.136
					(0.03)***
Dummy for New Sample					-34.844
					(11.012)**
Controls for Card Denomination/Brand	N	Y	N	N	N
Observations	10283	10283	8451	1832	10283
Mean of Dependent Variable	14.63	14.63	10.75	34.22	14.63
Sample	Both	Both	New	Old	Both

Notes: report (annualized) regression coefficients from regression of daily interest rate on these independent variables. Std. errors calculated by delta method.
p-value for test that Days Since first Visit at Shop + Days Since First Visit at Shop
* New Sample <0.001

Table 5: Correlations with Annual Return to Capital from Stock-out Sample using Variance Weighted Least Squares

	Independent Variable Characteristics		(1) Correlation with annual returns			(3) Correlation with log (annual returns)			Obs
	Mean	SE	Coeff	SE	t-stat	Coeff	SE	t-stat	
Years Education	10.9204	(3.221)	0.0034	(0.007)	0.475	0.0012	(0.001)	1.036	113
Years Shop Open	6.4596	(6.483)	0.0021	(0.006)	0.372	0.0008	(0.001)	0.854	111
Always keeps financial records	0.4602	(0.501)	-0.0311	(0.051)	0.610	-0.0119	(0.008)	1.430	113
Never keeps financial records	0.2920	(0.457)	-0.0025	(0.057)	0.044	-0.0001	(0.009)	0.015	113
Sometimes keeps financial records	0.2478	(0.434)	0.0434	(0.058)	0.744	0.0157	(0.009)	1.658	113
Always keeps purchase records	0.7679	(0.424)	-0.0197	(0.062)	0.319	-0.0075	(0.010)	0.747	112
Never keeps purchase records	0.0893	(0.286)	-0.0078	(0.086)	0.090	-0.0020	(0.014)	0.141	112
Sometimes keeps purchase records	0.1429	(0.351)	0.0377	(0.078)	0.485	0.0136	(0.013)	1.073	112
Other income last week	21.7439	(83.991)	0.0001	(0.001)	0.069	0.0000	(0.000)	0.155	113
Value of durable goods owned	2627.1500	(5974.907)	0.0000	(0.000)	1.563	0.0000	(0.000)	4.427	** 112
Value of animals owned	814.0836	(1063.893)	0.0000	(0.000)	0.978	0.0000	(0.000)	2.168	106
Got bank loan in past year	0.0973	(0.298)	0.0242	(0.078)	0.312	0.0093	(0.013)	0.729	113
Got MF loan in past year	0.1947	(0.398)	0.0027	(0.061)	0.045	-0.0010	(0.010)	0.098	113
has bank account	0.6637	(0.475)	0.0418	(0.055)	0.758	0.0142	(0.009)	1.587	113
ROSCA contributions in past year	95.0140	(164.350)	-0.0001	(0.000)	0.479	0.0000	(0.000)	1.117	112
If needed 1,000Ksh, take out of inventory	0.2743	(0.448)	0.0306	(0.055)	0.555	0.0109	(0.009)	1.212	113
If needed 10,000Ksh, take out of inventory	0.2124	(0.411)	0.0090	(0.057)	0.158	0.0035	(0.009)	0.379	113
Anything stolen in past year	0.0885	(0.285)	-0.0200	(0.112)	0.179	-0.0070	(0.018)	0.379	113
Any cards stolen in past year	0.0531	(0.225)	-0.0971	(0.366)	0.265	-0.0426	(0.060)	0.710	113
Value of cards stolen in past year	0.2831	(2.942)	0.1397	(0.147)	0.953	0.0297	(0.024)	1.214	108
Estimated monthly revenue	1952.2780	(5720.667)	0.0000	(0.000)	1.128	0.0000	(0.000)	2.977	** 94
Always gets bulk discount	0.2364	(0.427)	0.0547	(0.056)	0.970	0.0197	(0.009)	2.233	* 110
Sometimes gets bulk discount	0.2727	(0.447)	0.0058	(0.062)	0.093	0.0004	(0.010)	0.037	110
never get bulk discount	0.4909	(0.502)	-0.0535	(0.053)	1.001	-0.0180	(0.008)	2.147	* 110
Don't get the discount b/c lack the cash	0.7024	(0.460)	-0.0119	(0.068)	0.176	-0.0030	(0.011)	0.265	84
Don't get the discount b/c can make more investing elsewhere	0.0833	(0.278)	-0.7352	(0.275)	2.676	** -0.2986	(0.035)	8.455	** 84

if you had 1,000 Ksh, would invest in more inventory	0.7456	(0.437)	-0.0445	(0.060)	0.745	-0.0158	(0.009)	1.670	114
if you had 1,000 Ksh, would invest in bulk discounts	0.0351	(0.185)	0.0878	(0.098)	0.897	0.0311	(0.015)	2.009	114
if you had 1,000 ksh, would invest in other aspect of retail business	0.0175	(0.132)	123.6045	(53.753)	2.299 *	3.7625	(8.805)	0.427	114
if you had 1,000 ksh, would put in ROSCA	0.0439	(0.206)	0.1110	(0.112)	0.988	0.0384	(0.018)	2.167	114
if you had 1,000 ksh, would put in bank	0.0614	(0.241)	3.8083	(3.359)	1.134	0.6592	(0.538)	1.226	114
if you had 1,000 ksh, would invest in agriculture	0.0965	(0.297)	-0.0376	(0.081)	0.467	-0.0122	(0.013)	0.951	114
total startup costs	897.6158	(1470.029)	0.0000	(0.000)	0.085	0.0000	(0.000)	0.321	115
building startup costs	81.1548	(261.364)	-0.0001	(0.000)	0.211	0.0000	(0.000)	0.521	119
inventory startup costs	724.7650	(1432.807)	0.0000	(0.000)	0.036	0.0000	(0.000)	0.120	116
license startup costs	35.0366	(21.447)	-0.0007	(0.002)	0.379	-0.0004	(0.000)	1.174	118
other startup costs	55.4684	(95.443)	0.0000	(0.000)	0.091	0.0000	(0.000)	0.107	121

Notes: report (annualized) regression coefficients from regression of daily interest rate on these independent variables.