

PAYING MORE FOR LESS:
WHY DON'T HOUSEHOLDS IN TANZANIA TAKE ADVANTAGE OF
BULK DISCOUNTS?*

Brian Dillon[†]
Joachim De Weerd[‡]
Ted O'Donoghue[§]

October 4, 2016

Abstract

Despite average per-capita consumption of roughly \$1 per day, many Tanzanian households do not take advantage of bulk discounts when buying staple goods. Using transaction diaries covering nearly 57,000 purchases of 22 items by 1,499 households, we quantify the financial loss incurred (and the quantity of consumption forgone) by not taking advantage of bulk discounts that are available at frequently realized quantities. Using a conservative approach, we estimate that over two weeks the average household could spend 8.9% less on observed quantities (or consume 15.6% more at observed expenditure). We investigate several possible explanations for not buying in bulk, and find evidence consistent with inattention, worries about over-consumption, avoidance of social taxation, and household coordination problems. Contrary to prior work, we find little evidence that liquidity constraints prevent poor households from bulk purchasing.

JEL codes: O12, D03, D12

Keywords: bulk discounts; liquidity constraints; inattention; social taxes; coordination costs; self-control problems; consumer behavior

*For helpful discussions and comments we thank Jenny Aker, Chris Barrett, Jim Berry, Ann Bostrom, Peter Brummund, Arun Chandrasekhar, Paul Christian, Kelly Husted, Joe Kaboski, Kelsey Jack, Steve Kosack, Michael Kremer, Mujobu Moyo, Sendhil Mullainathan, Rohini Pande, Robert Plotnick, Jonathan Robinson, Mark Rosenzweig, Thaddeus Rweyemamu, Hilary Wething, Dean Yang, and seminar audiences at the ASSA conference, the CSAE conference in Oxford, University of Washington, University of Alabama, Cornell University, Michigan State University, and KU Leuven. Any errors are our responsibility.

[†]Corresponding author. University of Washington, Evans School of Public Policy and Governance, Box 353055, Seattle, WA, 98195. Telephone: (206)221-4601. Email: bdillon2@uw.edu.

[‡]IOB, University of Antwerp and LICOS, KU Leuven. Email: JoachimDeWeerd@UAntwerpen.be.

[§]Cornell University, Department of Economics. Email: edo1@cornell.edu.

1 Introduction

In Tanzania, average consumption per-capita is in the range of \$1 per day. At such low levels, the marginal value of additional consumption is very high. For this reason, we expect households in Tanzania and in other low-income countries to be especially mindful of opportunities to raise consumption through careful management of purchasing behavior.

In this paper we study the surprising, contrary finding that many households purchase non-perishable goods in small increments, multiple times, over a two-week period. If price schedules were linear and transaction costs minimal, frequent purchasing in small increments would have no impact on the budget set. We show, however, that many goods exhibit bulk discounts over commonly realized values of the quantity support. As a result, many households suffer significant financial losses (or, equivalently, substantial reductions in their consumption) in order to maintain a pattern of frequent, small-quantity purchases.

We have two main goals in the paper. First, we carefully document the existence of bulk discounts in markets in Tanzania, and quantify the financial losses that households incur (or the quantity of consumption forgone) by not taking advantage of these bulk discounts. Second, we empirically investigate why households do not take advantage of bulk discounts. Here, we focus on both traditional explanations—such as liquidity constraints, or fear of social taxation—and more behavioral explanations—such as inattention, or self-control problems.

In Section 2, we motivate our analysis with a simple example that illustrates how making frequent purchases in small quantities is financially inefficient when there are bulk discounts. We further use this example to delineate a variety of explanations for why a household might still choose to make frequent small-quantity purchases, despite incurring these losses. That discussion previews the possible explanations that we formally investigate later in the paper. We then develop our empirical framework. In particular, we describe how we use observed purchases at focal quantities—frequently observed quantities—to estimate a market expenditure schedule, $e_i^*(q)$, for every item i . This schedule returns the amount

one must spend to purchase quantity q of item i in a single transaction. The financial cost of small-quantity purchasing comes from comparing observed expenditure over two weeks to the estimated cost of purchasing the total observed quantity all in one purchase (using $e_i^*(q)$). In making these calculations, we project each observed transaction onto the market expenditure schedule, to ensure that the measure of financial losses is driven entirely by a failure to take advantage of bulk discounts, not by whether a particular transaction is a good or bad deal conditional on the quantity purchased.

We describe our data set in Section 3. The data come from the Survey of Household Welfare and Labour in Tanzania (SHWALITA), a consumption survey conducted in 2007-2008. We use a subset of the SHWALITA data, consisting of transaction diaries maintained by 1,499 households from 168 villages in 7 districts. For 2 weeks, these households recorded every purchase, gift, sale, or change in stocks, for all goods. We focus on the nearly 57,000 recorded purchases of 22 non-perishable items. We use the text description of each transaction to ensure that the item definitions are highly standardized. The set of items in the study includes numerous staple goods, such as kerosene, cooking oil, maize, flour, cassava, beans, rice, onions, cooking bananas, and dried sardines. Frequent purchasing is commonplace—among households that purchase an item, the average number of purchase transactions is 3.6 over two weeks.

In Section 4, we turn to our first main goal: carefully documenting the existence of bulk discounts, and quantifying the financial losses (or forgone consumption) from not taking advantage of these bulk discounts. Using focal quantities and prices to estimate $e_i^*(q)$ at the district level, we find that bulk discounts exist for 82 of the 146 item-district pairs. Some items exhibit bulk discounts in all seven districts, others in a subset of the districts. One item, brewed tea, does not exhibit bulk discounts in any district.¹

¹There is a longstanding literature on the estimation of demand elasticities for consumer goods in developing countries (Deaton, 1988; Subramanian and Deaton, 1996; Deaton, Friedman and Alatas, 2004; McKelvey, 2011). That literature focuses on the problem of using unit values from aggregate data to proxy for prices, because transaction-level data like ours is rarely available. Those papers do not allow for bulk discounts, as identification usually requires assuming linear prices at the local level.

We find that households appear to be sacrificing a substantial amount of consumption by not taking advantage of bulk discounts. Across items purchased multiple times in the two-week study period, the value of forgone consumption is equal to 8.9% of the value of expenditure. In other words, the average household could spend nearly 9% less on a range of important goods without reducing consumption. This average value masks significant heterogeneity. Approximately 8% of households have zero forgone consumption, while a quarter could reduce expenditure by over 10% without reducing quantity purchased. If we take the alternative approach of holding expenditure constant and calculating the counterfactual quantities that could be purchased by buying all at once, we find ever larger average values. Households could purchase 33% more kerosene, 50% more cooking bananas, 26% more cooking oil, 46% more onions, 28% more sardines – surprisingly large amounts for goods that are part of daily life in Tanzanian villages.

In Section 5, we address our second main goal: investigating why households fail to take advantage of bulk discounts. Recent papers have shown how seemingly inexplicable behavior by individuals in developing countries can be understood by either gathering better information about the setting (Baland, Guirkinger and Mali, 2011; Stephens and Barrett, 2011; Zeitlin, 2011; Burke, 2014), or expanding the class of admissible explanations to include those grounded in models that allow for psychological biases or cognitive costs (Ashraf, Karlan and Yin, 2006; Duflo, Kremer and Robinson, 2011; Dupas and Robinson, 2013; Mani et al., 2013; Kremer et al., 2013; Jack and Smith, 2015; Kremer et al., 2015). We borrow from both of these approaches, and consider a wide range of both rational and behavioral mechanisms.

We first focus on explanations that we conclude are not playing a major role in driving the losses that we estimate. The hypotheses that we reject are the following: binding liquidity constraints prevent households from buying in bulk; it is costly to store or transport bulk quantities; consumers are not aware of bulk discounts; people enjoy going to the market and shopping; and, paying high mark-ups on small quantities of some items is important for

maintaining buyer-seller relationships. Evidence against these hypotheses comes primarily from the rejection of necessary conditions, through a combination of parametric and non-parametric arguments. Our finding that losses are not driven by liquidity constraints is particularly important relative to the prior literature on this question, which has largely found that an inability to take advantage of bulk discounts leads the poor to pay more for consumption (Rao, 2000; Attanasio and Frayne, 2006; Gibson and Kim, 2011; Mussa, 2015; Attanasio and Pastorino, 2015). In contrast, we find that losses from not buying in bulk are not concentrated among poorer households. Moreover, it is not the case that poorer households tend to make purchases in smaller quantities than wealthier households.

Next, we focus on explanations for which we find supportive evidence. First, while we conclude that households are surely aware of the available bulk discounts, loss-prone households seem to be inattentive to the financial implications of not buying in bulk. Accounting for both household and item fixed effects, loss-prone households do not reduce the number of purchase transactions on items with bulk discounts, whereas other households do. Second, we find evidence that households make more frequent purchases as a way of rationing their consumption, perhaps due to worries that they will consume large stocks more quickly than they would like. Evidence related to rationing is based on the analysis of shopping patterns for “temptation goods”, which we identified through a separate survey that we conducted in Tanzania. Third, we find evidence that households might purchase in small quantities in order to avoid social taxation—in the form of friends, family, and neighbors consuming a portion of their stocks (Anderson and Baland, 2002; Platteau, 2006; Goldberg, 2016; Baland, Guirkingier and Mali, 2011; Alby, Auriol and Nguimkeu, 2013; Jakiela and Ozier, 2016). In fact, because we observe flows of both incoming and outgoing resources, we can construct a proxy measure of each household’s social tax rate. Finally, we find evidence that losses are larger in households that seem to have a harder time coordinating their purchases, as proxied by the number of times that two members of the same household purchase the same item on the same day. This finding suggests that each household member makes purchases

in small quantities as a kind of hedge against the possibility that another household member is purchasing the same item.

Finally, in Section 6, we discuss some broader implications of our analysis. We highlight the that household purchasing behavior seems driven by multiple mechanisms. Moreover, these mechanisms might interact with each other in important ways, and operate differently across different households. We also discuss possible policy implications.

2 Conceptual framework

In this section, we develop a conceptual framework that underlies our empirical analysis. We begin with a stylized, motivating example of bulk discounts. We then build upon this example to introduce our approach to the data.

2.1 Motivating example

Suppose that, in a particular district of Tanzania, rice is sold in three quantities: 1 kg for 1000 TZS; 2 kg for 1600 TZS; and 4 kg for 2400 TZS. Note the existence of bulk discounts: the unit (per-kg) price is 1000 TZS for the 1 kg purchase, 800 TZS for 2 kg, and 600 TZS for 4 kg.

A household that wishes to consume 4 kg of rice over a two-week period has (at least) three options: it could purchase a 4 kg bag and consume it over the two weeks; it could purchase a 2 kg bag, consume it over the first week, then purchase another 2 kg bag at the start of the second week; or it could purchase a 1 kg bag four times over the course of the two weeks. From a purely financial perspective, purchasing the 4 kg in a single transaction—buying in bulk—is most efficient. Yet, we will see that many households do not always purchase in bulk.

Suppose we observe a household that purchases a 1 kg bag on four occasions over the two weeks. This raises two questions. First, how large are the losses incurred from not

buying in bulk? Second, why might this household be willing to incur the loss?

In our stylized example, the answer to the first question is straightforward. We can calculate a financial loss by taking what the household actually paid for the 4 kg of rice (4000 TZS) and subtracting what the household could have paid had it purchased the entire 4 kg in bulk (2400 TZS). Hence, this household incurred a financial loss of 1600 TZS for rice. Alternatively, we can calculate a quantity loss by taking the total amount spent (4000 TZS), calculating the quantity that could have been purchased had that entire amount been spent at the lowest per-unit price ($6\frac{2}{3}$ kg at 600 TZS per kg), and subtracting the quantity the household actually purchased (4 kg). Hence, this household incurred a quantity loss of $2\frac{2}{3}$ kg. With some minor adjustments, this is the approach that we take in the data (see the next subsection for details).

Turning to the second question, there are a number of reasons a household might choose not to buy in bulk despite the financial cost entailed. To foreshadow our analysis in Section 5, here we provide a brief overview of the hypotheses we will consider. These hypotheses arose from three sources: discussions with individuals and focus groups in Tanzania (described in Section 3), our own hypothesizing based on the literature or our knowledge of the context, and suggestions from colleagues or seminar participants.

The first set of possible factors relates to the full cost of getting goods from the market to the house. A natural concern in Tanzania is that the household may be liquidity constrained. If household members are living on day-to-day wage earnings, have minimal savings, and must satisfy a range of consumption needs, they may not be able to spend 2400 TZS on rice in a single purchase. This theme is prominent in existing work on bulk discounts. Another possibility is that there are other costs associated with bulk purchasing, e.g. transport costs, storage costs, or expected losses to depreciation. A full accounting of such costs might show that the bulk purchase is not optimal.

A second set of factors revolves around the way that goods are consumed. Prior evidence suggests that in sub-Saharan Africa, wealth is subject to a high rate of “social

taxation” in the form of friends and relatives asking for handouts or coming over for meals. Pressure to distribute resources among social networks may be particularly significant for visible goods such stores of staple items. In our example, if the household buys in bulk, friends and relatives may notice the large bag of rice and be more inclined to ask for a loan or to eat a meal at the house. Relatedly, a household might be worried about self-control problems among its members. If having 4 kg of rice makes household members more prone to consume a little extra rice on any given day, the household could end up consuming its rice store more quickly than planned.

A final set of possible factors relates to how people shop. Markets are gathering places and centers of social life. There could be utility value from visiting the market frequently, and making regular purchases may be a part of the implicit social contract. Or, it may be the case that bulk discounts are a function of the relationships between buyers and sellers, so that discounts are only available to those who make regular purchases from across the quantity distribution. Additionally, it may be costly to coordinate purchases within the household. For instance, if a person is at the market and considers buying some rice, she may wonder whether someone else from the household has already purchased rice. To avoid both a shortage and an excess of rice at home, it might be natural to purchase only 1 or 2 kg. Finally, it may be the case that when making purchases, shoppers are simply inattentive to the magnitude of forgone consumption from purchasing in small quantities.

2.2 Framework for empirical analysis

In order to quantify the losses from not buying in bulk, we generalize the approach in our stylized example to reflect nuances in the data. We again begin with a set of focal quantities: suppose an item i is available in the market in R focal quantities $\{q_r\}_{r=1}^R$, ordered so that $q_1 < q_2 < \dots < q_R$. Let e_r denote the expenditure required to purchase quantity q_r , and let p_r denote the associated unit price, so $p_r = e_r/q_r$. Suppose the focal quantities (weakly) exhibit bulk discounts, so that $p_1 \geq p_2 \geq \dots \geq p_R$. A focal quantity should be interpreted as

roughly a package size or a common unit of trade, analogous to the three quantities in our rice example. In some cases these focal quantities correspond to actual package sizes from mass produced items, such as 1 liter bottles of cooking oil. In other cases, local units have emerged over time as vendors adopt widely available buckets or canisters as standard units of trade.

Our approach will be to identify focal quantities empirically, using commonly observed transaction quantities. In the following section we will be precise about how we do this, and about how we estimate focal prices. For the moment, take it as given that we observe $\{q_r\}_{r=1}^R$ and $\{p_r\}_{r=1}^R$ for each item.

Over the two weeks during which it is observed, suppose that household h buys item i in K separate transactions. Let $k = 1, \dots, K$ index the household's purchases, with the associated quantities and expenditures denoted q_{hik} and e_{hik} , respectively. Household h 's *observed total expenditure* on item i is $e_{hi} \equiv \sum_{k=1}^K e_{hik}$, and its *observed total quantity* of item i is $q_{hi} \equiv \sum_{k=1}^K q_{hik}$. As in our motivating example, our goal is to calculate (i) the financial savings if household h had instead purchased its entire observed total quantity in a single transaction, and (ii) the extra quantity if household h had instead spent its entire observed total expenditure in a single transaction. Before we can do so, however, we must address some issues that arise in the data.

First, these calculations require knowing the expenditure associated with any quantity, not just focal quantities. To reflect the reality of shopping in these markets, we base such estimates on the expenditure required to purchase a particular quantity using only focal quantities. Specifically, we define the *expenditure schedule*, $e_i^*(q)$, as follows. For any focal quantity q_r , we assign $e_i^*(q_r) \equiv e_r$. For any quantity that is between two focal quantities, we use the weighted average of the expenditures for the closest focal quantities on either side:

$$\text{for any } q \in [q_r, q_{r+1}], e_i^*(q) \equiv \left(\frac{q_{r+1} - q}{q_{r+1} - q_r} \right) e_r + \left(\frac{q - q_r}{q_{r+1} - q_r} \right) e_{r+1}.$$

Finally, we assign the lowest unit price (p_R) to any quantity greater than the largest focal

quantity (i.e., $e_i^*(q) \equiv qp_R$ for any $q > q_R$), and the highest unit price (p_1) to any quantity less than the smallest focal quantity (i.e., $e_i^*(q) \equiv qp_1$ for any $q < q_1$). Note that this expenditure schedule can be converted into a *unit price schedule* using $p_i^*(q) = e_i^*(q)/q$.

There are two interpretations of this weighted average approach. The first relates to actual behavior in the market. Consider a shopper in our motivating example trying to buy 3 kg of rice in a single transaction. The buyer may argue that she should pay at most the 2 kg unit price, and perhaps an even lower unit price, given how much rice she is buying. If the probability of receiving a particular unit price is proportional to the difference between the quantity being purchased and the nearest focal quantities, our measure assigns the expected value. A second interpretation relates to our choice of an aggregation period of two weeks, which is necessarily somewhat arbitrary. In most cases, q_{hi} , the aggregate quantity purchased over two weeks, will not correspond to an exact focal quantity. But we could just as easily aggregate purchases over a longer or shorter time period to ensure that q_{hi} is equal to a focal quantity. We are effectively calculating the expenditure associated with aggregating to the next lower or next higher focal quantity, and then taking a weighted average.

A second issue is that while the majority of observed transactions take place at focal quantities, some do not. Moreover, some transactions at focal quantities are not at focal prices. To deal with idiosyncratic variation in quantities and unit prices, we project all observed transactions onto the expenditure schedule prior to aggregation. Hence, observed expenditure by household h on item i in transaction k can be represented as $e_{hik} = e_i^*(q_{hik}) + \nu_{hik}$, where ν_{hik} is an idiosyncratic component. Then, household h 's *adjusted expenditure* on item i in transaction k is $\hat{e}_{hik} \equiv e_i^*(q_{hik})$, and its *adjusted total expenditure* on item i is $\hat{e}_{hi} \equiv \sum_{k=1}^K \hat{e}_{hik}$. Using adjusted total expenditures in our calculations of losses will ensure that our results are not distorted when a household's actual expenditure in a particular transaction happens to be above or below the expenditure schedule.

Although we will show that bulk discounts are clearly identifiable within-village and even within-household, for power reasons we will construct $e_i^*(q)$ at the district level. Yet,

households within a district might face slightly different prices due to variation in bargaining power, between-village price variation, and residual variation in unobserved quality (in Section 3 we explain how we reduce the possibility of unobserved quality variation). However, because we use adjusted rather than observed expenditures, we do not need to assume that every household in a district faces exactly the same schedule. Instead, the identifying assumption is that any differences in expenditure schedules take the form of parallel shifts over the relevant ranges. This will not be especially restrictive, because the ranges of quantities under consideration here are infra-marginal to wholesale volumes, so that the transport cost component of any between-village price variation does not vary across focal points.

With this framework, we can calculate the financial savings to household h if it had purchased its entire observed total quantity over the two weeks in a single transaction. Specifically, household h 's *financial loss* on item i is $L_{hi} \equiv \hat{e}_{hi} - e_i^*(q_{hi}) = \left(\sum_{k=1}^K e_i^*(q_{hik}) \right) - e_i^*(q_{hi})$, and its *percentage financial loss* is $\tilde{L}_{hi} \equiv L_{hi}/\hat{e}_{hi}$. By construction, L_{hi} and \tilde{L}_{hi} are zero for any item that the household purchases only once over the two weeks. This makes our approach conservative, because the items we study are popular consumer goods in Tanzania and in many cases they can be stored for months. Households that purchase item i only once over the study period, in small quantity, could in all likelihood reduce expenditure by bulk purchasing for a longer time period.

Alternatively, we can also calculate the extra quantity that household h could purchase if it spent its total adjusted expenditure in a single transaction. Specifically, inverting the expenditure schedule, household h 's *quantity loss* on item i is $Q_{hi} \equiv e_i^{*-1}(\hat{e}_{hi}) - q_{hi}$, and its *percentage quantity loss* is $\tilde{Q}_{hi} \equiv Q_{hi}/q_{hi}$. Again, by construction, Q_{hi} and \tilde{Q}_{hi} are zero for any item that a household purchases only once over the two weeks. For much of the analysis, we will focus on the financial loss measures, L_{hi} and \tilde{L}_{hi} , because they can easily be aggregated across items (we will typically refer to these as “loss” and “percentage loss”). However, the quantity loss measures provide an additional way to understand the magnitude of the purchasing inefficiencies in the data.




Figure I: Expenditure and unit price for the example with 3 focal points
Notes: Authors' calculations from example data in text.

Figure I gives an example our approach. Imagine a household that buys rice in the market described in Section 2.1. The household reports three rice purchases over the observation period: 1 kg for 1000 TZS, 1 kg for 1250 TZS, and 1.5 kg for 1500 TZS. The \times in Figure I mark the actual transactions, with the unit price schedule in the left panel and the expenditure schedule in the right panel. Observed expenditure is 3750 TZS (point A). Adjusted expenditure is $e^*(1) + e^*(1) + e^*(1.5) = 1000 + 1000 + 1300 = 3300$ (point B). Counterfactual expenditure from bulk purchasing is $e^*(3.5) = 2200$ (point C). These three expenditure values are associated with the total observed quantity of 3.5 kg. Finally, the counterfactual quantity that could be purchased using the total adjusted expenditure of 3300 all at once is $e_i^{*-1}(3300) = 3300/600 = 5.5$ kg (point D). For this example, the financial measures of loss are $L_{hi} = 3300 - 2200 = 1100$ (the vertical distance between points B and C), and $\tilde{L}_{hi} = 1100/3300 = 33.3\%$. The quantity measures of loss are $Q_{hi} = 5.5 - 3.5 = 2$ kg (the horizontal distance between points B and D), and $\tilde{Q}_{hi} = 2/3.5 = 57.1\%$. This household could reduce its expenditure on rice by 33.3% without reducing its quantity consumed, or it could increase its quantity of rice consumed by 57.1% without increasing its expenditure.

3 Data and descriptive patterns

The data for this paper are from the Survey of Household Welfare and Labour in Tanzania (SHWALITA). The survey was part of an experiment to test the impact of questionnaire design on consumption measures (see Beegle et al. (2012) for details). In one arm of the study, 9 households per village were randomly assigned to complete a consumption diary in one of three conditions. Three households completed a single, household-level diary, with no monitoring by project staff. Three households also completed a single, household-level diary, but received multiple follow-up visits from an enumerator or local assistant. Lastly, in the final three households, each adult member kept his or her own diary, with children placed on the diaries of the adults who knew most about their daily activities. Households in the third group received multiple follow-up visits, similar to those in the second group. The differences between module arms are small but non-negligible, and they have no impact on the findings in this paper. We control for the diary module type in all relevant regressions.

The SHWALITA survey was conducted in 24 villages per district, in 7 districts. The resulting data set includes responses for 1,512 diary households. We dropped 10 households that did not purchase any of the items that we study, and 3 that did not complete the end-line survey, leaving a sample size of 1,499. Data were collected from September 2007 to August 2008. All households in a given village completed their diaries over the same 14 days. Survey work in each district was completed in less than two months.²

Each day during the study period, diary keepers recorded the quantity, unit, value, and description of all items that came into or went out of the household. Respondents separately listed purchases, gifts, own production, and stock adjustments. We use only the purchasing data, except where otherwise specified. After the diary period, project staff assigned each listed item to one of 73 categories, covering 58 food items and 15 non-food items. For this paper we drop items with too few observations, and drop perishable items

²More details are available at the project page, accessible here: <http://edi-global.com/publications/>. Data are available by contacting the lead SHWALITA researchers listed on the project page.

that cannot be stored for two weeks by most households (such as beef, milk, and fresh fish). We include one service – milling of staple grains – which is frequently purchased. Importantly, we do not drop or retain items based on whether the price schedule exhibits bulk discounts. This improves the generalizability of our findings, and is useful for some of the tests implemented in Section 5.

To minimize quality variation within items, we use the detailed transaction descriptions to standardize items. For example, we drop “unrefined sugar” from “Sugar”, retain only “dried beans” from the original category of “Peas, beans, lentils and other pulses”, keep only “immature coconuts” in the “Coconut” category, and restrict the “Dried fish” group to only “dried sardines”, excluding larger fish. In this way we create 22 items that are far more uniform than the goods in a typical consumption survey. Table I shows details.

A further cleaning step was required to standardize units. Respondents reported quantities in kilograms, liters, and a range of local units such as bunches, heaps, tins, ladles, buckets or bundles. Field team members conducted a market price survey in each village in conjunction with the household survey, and made a point of precisely measuring local units at the market. Staff members recorded the most common units in which each item was sold, measured the unit in kilograms or liters, and noted the price. This was done for three vendors per market, and at multiple markets if there was more than one major market in a village. If a diary item was not sufficiently covered by the market price survey, the survey team purchased and weighed the item in question. We use the median, district-level conversion rates from these surveys to convert local units into kilograms or liters.

The final data set contains details for 56,892 separate transactions. From the purchase quantity statistics in Table I we see that maize and cooking bananas, staple carbohydrates, are the items purchased in the largest kilogram quantities. Comparing the average transaction quantities to the average total 2-week purchase quantities, it is clear that the average household buys the items it consumes multiple times over two weeks.

Table II shows the pattern of purchases and expenditures, across households, over the

Table I: Descriptions of item standardization and units

Description of item	Avg. transaction quantity	Avg. total 2 week quantity	Unit
Maize: loose, dried maize kernels. Excludes maize flour, maize cobs, popcorn, or processed maize grains.	10.36	20.90	Kg
Milling: fee paid for machine-grinding. Mostly maize, but milling of millet, sorghum and rice is not excluded. Husking rice is excluded.	8.60	20.57	Kg
Cooking Bananas: excludes any other type of banana such as roasting bananas, beer bananas or sweet bananas.	7.63	17.32	Kg
Cassava: fresh, raw cassava. Excludes cassava flour and dried, boiled, fried, or roasted cassava.	3.18	7.52	Kg
Soap: solid bar soap. Excludes powdered soap, beauty soap, dishwashing liquid.	2.21	6.56	Kg
Charcoal: excludes wood, kerosene, other fuels for cooking.	2.06	13.05	Kg
Rice: husked white rice. Excludes unhusked, brown, broken rice.	1.63	5.86	Kg
Flour: white maize flour. Excludes brown flour, flours from wheat, millet, sorghum.	1.27	7.19	Kg
Beans: dried kidney beans. Excludes fresh kidney beans, green beans, other beans, green gram, lentils, chick peas, cow peas, pigeon peas, bambarra nuts, garden peas.	0.85	2.19	Kg
Coconut: whole matured coconuts. Excludes immature coconuts.	0.76	3.88	Kg
Tomatoes: fresh, whole tomatoes. Excludes cherry tomatoes and canned tomatoes.	0.57	2.85	Kg
Sugar: refined sugar. Excludes unrefined sugar, honey, syrup, other sweeteners.	0.54	2.22	Kg
Salt: excludes coarse salt or any other spices.	0.52	1.06	Kg
Sweet Bananas: excludes cooking, roasting or beer bananas.	0.42	0.97	Kg
Sardines dried <i>dagaa</i> . Excludes fresh <i>dagaa</i> and other fish.	0.36	1.22	Kg
Onions: fresh, whole onions.	0.30	1.15	Kg
Tea Leaves: black tea leaves. Excludes other types of tea, ground coffee, instant coffee and other raw ingredients for hot beverages.	0.02	0.08	Kg
Tea: black tea with milk, ready to drink. Exclude other types of tea, coffee, cocoa, and any other hot beverage.	0.75	3.32	Liter
Kerosene: very homogenous so no need to exclude anything in this category. Typically used for lighting and/or cooking.	0.26	0.95	Liter
Cooking Oil: liquid vegetable oil. Excludes, butter, ghee, other types of fat.	0.19	0.84	Liter
Cigarettes: Portsman cigarettes. Excludes other brands, locally made cigarettes, chewing tobacco, and raw tobacco.	5.10	30.69	Piece
Matches: excludes lighters or wicks.	1.97	4.18	Box

Notes: Authors' calculations from SHWALITA data.

two-week diary period. The total number of observed transactions ranges from 688 (maize) to 5472 (tomatoes). The average item was purchased by just under half of the sample (725 households), and was purchased multiple times by just over a third of the sample. Some items, such as sugar, tomatoes, sardines, onions, cooking oil, and kerosene, were purchased

Table II: Purchase and expenditure patterns, by item, N=1499 households

Item	Total purchases	HHs purchasing	HHs multiple purchasing	Among all households			Among households purchasing item		
				% of HHs multiple purchasing	Avg no. of purchases	Avg total expenditure	% of HHs multiple purchasing	Avg no. of purchases	Avg total expenditure
Tomatoes	5472	1101	893	59.6	3.7	870	81.1	5.0	1184
Cooking Oil	5297	1164	925	61.7	3.5	1498	79.5	4.6	1929
Kerosene	4558	1271	976	65.1	3.0	1305	76.8	3.6	1540
Sugar	4293	1044	789	52.6	2.9	1832	75.6	4.1	2631
Onions	4127	1089	799	53.3	2.8	370	73.4	3.8	509
Flour	4042	715	556	37.1	2.7	2051	77.8	5.7	4300
Sardines	3523	1055	796	53.1	2.4	800	75.5	3.3	1136
Rice	3303	919	631	42.1	2.2	3070	68.7	3.6	5007
Soap	3202	1078	747	49.8	2.1	651	69.3	3.0	906
Salt	2299	1125	686	45.8	1.5	427	61.0	2.0	569
Tea Leaves	2227	611	391	26.1	1.5	157	64.0	3.6	385
Beans	2051	797	501	33.4	1.4	1154	62.9	2.6	2172
Matches	1958	922	522	34.8	1.3	119	56.6	2.1	194
Coconut	1905	373	312	20.8	1.3	416	83.6	5.1	1672
Charcoal	1713	270	230	15.3	1.1	892	85.2	6.3	4955
Tea	1565	351	236	15.7	1.0	164	67.2	4.5	703
Milling	1356	567	309	20.6	0.9	301	54.5	2.4	796
Cigarettes	957	159	116	7.7	0.6	162	73.0	6.0	1532
Cassava	881	372	192	12.8	0.6	245	51.6	2.4	988
Sweet Bananas	749	325	145	9.7	0.5	162	44.6	2.3	749
Cooking Bananas	726	320	173	11.5	0.5	508	54.1	2.3	2383
Maize	688	341	143	9.5	0.5	1673	41.9	2.0	7354
AVERAGE	2586	725	503	33.6	1.7	856	67.2	3.6	1982

Notes: Authors' calculations from SHWALITA data.

Table III: Summary statistics at the household level

	Mean	s.d.	Median
Age of head (years)	46.66	16.03	44.00
Education of head (years)	4.73	3.75	7.00
Head is female (=1)	0.20	0.40	0.00
Household size	5.33	2.96	5.00
Share under 15 yrs old	0.42	0.24	0.50
Share over 65 yrs old	0.07	0.19	0.00
Urban area (=1)	0.34	0.48	0.00
Acres owned	3.83	5.56	2.00
Wealth index value	-0.01	1.00	-0.43
Nominal consumption (TZS/yr)	2001642	1974544	1449216
Nominal consumption (USD/yr)	1741	1717	1260
Nominal consumption per capita (TZS/yr)	450154	469498	304887
Nominal consumption per capita (USD/yr)	391	408	265

Notes: Authors' calculations from SHWALITA data. Sample size is 1,497, because two households with incomplete demographic data are not included. Tanzania shillings converted to USD at a rate of 1,150 TZS/USD.

more than once by a majority of households. Among only the households that purchase each item, the highest average expenditure is on maize at 7,354 TZS/household, and the lowest is on matches at 194 TZS/household. The average number of purchases per item is 3.6.

In Table III we report household summary statistics. Mean consumption per capita is almost 400 USD per year, but the distribution is heavily skewed; the median is only 265 USD per year.³ The median household has 5 people. The “Wealth index value” is the value of the first principal component from a vector of household assets (Filmer and Pritchett, 2001; Sahn and Stifel, 2003). The assets used to construct this wealth index include dwelling characteristics such as roof material, wall material, and number of rooms, as well as ownership of durable goods such as phones, other electronics, and bicycles. We use this index as our primary measure of household wealth because it is not endogenous to concurrent consumer prices, unlike consumption or expenditure.

While writing this paper we conducted two additional, small-scale data collection efforts. The first consisted of informal interviews and focus groups with people in Tanzania,

³When applicable, Tanzania shillings (TZS) are converted to US dollars (USD) at the rate of 1,150 TZS/\$1, the average exchange rate from the survey period.

during the years 2012-2015. We interviewed roughly 10 individuals, mostly consumers but also a few shopkeepers. We conducted three informal focus groups, two in the Kagera region and one in the Dodoma region, with 5-6 people at each. Through these interviews and discussions we identified some of the hypotheses for why households might forego bulk discounts, and heard stories about household shopping patterns that helped steer the analysis. As qualitative work goes, our efforts were decidedly informal. Yet, these conversations were highly informative. We make one mild empirical claim based on this work (in Section 5.2.3), and we refer to some of the comments from interviewees when relevant.

Our second data collection effort was an on-line survey conducted in June-July 2016. This short survey was sent to a group of Tanzanians with extensive experience studying household decision-making around economic issues. We asked these respondents to characterize items by the likelihood that they are “temptation goods”, i.e., items that one tends to over-consume (relative to a plan) when held in stock. We describe this survey in more detail in Section 5.2.2.

4 Quantifying the value of forgone consumption

In this section we estimate the value of consumption that households forego by not buying in bulk. In Section 4.1 we describe the bulk discounts and the estimated expenditure schedules. In Section 4.2 we estimate counterfactual expenditures, holding quantity fixed, and counterfactual quantity, holding expenditure fixed. We also examine the distributions and correlates of losses. Finally, in Section 4.3 we examine the variation around the expenditure schedule, to understand why households sometimes pay different prices for the same quantity.

4.1 Bulk discounts in the SHWALITA data

While our main analysis uses the focal price approach described in Section 2.2, it is instructive to first estimate linear unit price and expenditure schedules. Columns 1 and 2 of Table

IV show slope coefficients from a series of item-specific, transaction-level regressions, with unit price as the dependent variable and quantity as the independent variable, using the transaction diary data. Each coefficient is from a separate regression. Every estimated slope coefficient is negative. This holds both within village-day (column 1), where 13 of 22 slope coefficients are negative and significant, and within household (column 2), where 17 of 22 coefficients are negative and significant. The clear pattern is that unit price is decreasing in quantity for many items.

In columns 3–4 of Table IV we report the slope coefficients from regressions of unit price on quantity using the market price survey data. Survey team members collected these data directly at local markets. The underlying regressions include district (column 3) or village (column 4) fixed effects, with standard errors clustered at the district level. The market price data do not cover as many items as the diaries, and span a smaller range of quantities, leading to some level differences between the two sets of slope coefficients. However, the large majority of coefficients in columns 3 and 4 are negative and statistically significant. As with the diary data, there are no positive and significant slope coefficients. An important implication of these estimates is that bulk discounts are not just for regular customers, but are available to all shoppers, including survey team members who are not local.⁴

Having established that bulk discounts exist, we now pursue the non-parametric approach described in Section 2.2. We designate a quantity as focal if it accounts for at least 5% of all observations at the item-district level.⁵ By this definition there are 1-9 focal quantities

⁴One might be concerned about division bias from constructing unit price as the quotient of two variables measured with error. To examine this, we estimated expenditure schedules by regressing transaction-level expenditure on quantity and its square, suppressing the constant and fixed effects to enforce regression through the origin. The coefficient on q^2 is negative and significant for 10 of 22 items, negative and close to significant for 5 others, and never positive and significant. The implication is that expenditure schedules are generally concave, which is consistent with bulk discounts. Results in online appendix Table A.1.

⁵We ignore the roughly 1 in 5 candidate focal quantities that either require greater total expenditure than a larger-quantity focal point, or that have a higher unit price than a smaller-quantity focal point. These points can never be part of an optimal counterfactual purchase. Because this also impacts adjusted expenditure, dropping these points has the effect of substantially attenuating the estimated losses (average losses are nearly 70% higher if we do not drop these points). This is another way that our approach is conservative in estimating losses.

Table IV: Regressions of unit price on quantity, various specifications

Dependent variable: transaction-level unit price				
Item	Transaction diary data		Market surveys	
	(1)	(2)	(3)	(4)
Rice	-4.7* (2.0)	-9.5** (2.6)	-43.3* (19.1)	-26.7 (44.0)
Maize	-3.7* (1.6)	-6.1** (2.4)	-1.5*** (0.3)	-1.9 (1.1)
Flour	-9.6** (3.2)	-7.8** (3.1)	-171.2 (119.3)	-312.3 (161.3)
Milling	-0.9* (0.3)	-0.9** (0.3)		
Cassava	-3.9 (9.0)	-2.3 (1.9)	-25.6** (8.6)	-221.7*** (20.0)
Cooking Bananas	-40.7* (16.9)	-23.4** (6.3)	1.6 (5.5)	7.5 (7.8)
Sugar	-94.5* (40.0)	-78.1** (29.6)	-97.8*** (14.2)	-77.6*** (16.8)
Beans	-17.9 (24.0)	-32.9* (15.8)	-212.3* (97.8)	-198.9 (188.1)
Coconut	-21.6 (15.3)	-47.5** (9.4)		
Tomatoes	-83.1 (46.5)	-129.5* (63.0)		
Onions	-419.0** (119.9)	-396.4*** (79.0)		
Sweet Bananas	-369.6 (403.5)	-210.4* (89.3)	-32.9*** (8.0)	-32.1** (12.1)
Sardines	-371.5** (129.4)	-336.4** (127.4)	-149.1 (208.3)	-272.0 (260.8)
Cooking Oil	-1135.8*** (253.3)	-1199.8*** (239.5)	-756.3** (252.0)	-731.2* (315.6)
Salt	-358.9 (320.6)	-267.2 (215.6)		
Tea	-9.6 (15.9)	-8.6 (18.5)		
Tea Leaves	-27319.1** (7765.7)	-27338.0** (8420.1)		
Charcoal	-17.0** (3.9)	-19.8 (9.7)		
Kerosene	-1378.3** (425.0)	-1206.5*** (297.6)		
Matches	-2.5 (2.6)	-2.7 (2.3)		
Soap	-2.4* (1.2)	-4.7** (1.6)		
Cigarettes	-0.1 (0.1)	-0.1** (0.0)		
Fixed effects	Village-day	Household	District	Village

Notes: Authors' calculations from SHWALITA data. Standard errors in parentheses; standard errors clustered at district level; *** sig. at 0.01, ** sig. at 0.05, * sig. at 0.1. Controls for diary type included in regressions underlying columns 1 and 2. Each coefficient is from a separate regression of unit price on quantity, for only the item indicated.

per item-district pair, with just over 3.3 on average. Across all items and districts, 70.2% of purchases were made at a focal quantity. We use the median unit price at the focal quantity as an estimate of the focal price.

Table V: Example focal quantities and prices from a single district

Item	Statistic	Focal point					cumulative % coverage
		1	2	3	4	5	
Rice	Quantity	.5	1	1.5	2	3	
	Frequency (%)	12.3	42.6	10.5	16.4	6	87.8
	Median unit price	1000	1000	1000	1000	1000	
	Expenditure	500	1000	1500	2000	3000	
Sugar	Quantity	.05	.25	.5	1	.	
	Frequency (%)	6.2	50.1	17.6	17.2	.	91.1
	Median unit price	2000	1000	1000	1000	.	
	Expenditure	100	250	500	1000	.	
Tomatoes	Quantity	.35	1.05	.	.	.	
	Frequency (%)	41.8	6.1	.	.	.	47.9
	Median unit price	571	286	.	.	.	
	Expenditure	200	300	.	.	.	
Onions	Quantity	.05	.1	.35	.7	.	
	Frequency (%)	30.6	9.1	36.8	11.3	.	87.8
	Median unit price	1000	1000	286	286	.	
	Expenditure	50	100	100	200	.	
Sardines	Quantity	.1083	.2167	.325	.	.	
	Frequency (%)	12.8	17.6	6.6	.	.	37.0
	Median unit price	1385	923	923	.	.	
	Expenditure	150	200	300	.	.	
Cooking Oil	Quantity	.045	.09	.135	.5	.	
	Frequency (%)	29.1	30.2	5.2	5.1	.	69.6
	Median unit price	3333	3333	3333	2200	.	
	Expenditure	150	300	450	1100	.	
Kerosene	Quantity	.045	.09	.135	1	.	
	Frequency (%)	32.1	18.7	9.1	9.1	.	69.0
	Median unit price	2222	2222	1556	1200	.	
	Expenditure	100	200	210	1200	.	
Soap	Quantity	1	2	.	.	.	
	Frequency (%)	68.7	20.5	.	.	.	89.2
	Median unit price	150	150	.	.	.	
	Expenditure	150	300	.	.	.	

Notes: Authors' calculations from SHWALITA data. Focal points are quantities that account for at least 5% of transactions at the district level.

Table V shows descriptive statistics for a set of example price schedules from one of the study districts. Quantities are in kilogram units for all goods other than cooking oil and kerosene, which are in liters. Rows show the quantity, percentage of purchases, median unit price, and expenditure at each focal point. Bulk discounts are visible in the unit price

data: the minimum unit price becomes available at 0.25 kg of sugar, 0.5 liters of cooking oil, 0.35 kg of onions, and so on. For some of the tests in Section 5 it will be useful to define q_{min}^* , the minimum quantity that must be purchased in order to reach the lowest available unit price, and e_{min}^* , the minimum expenditure required to access the lowest available unit price (i.e., to purchase q_{min}^*). In Table V, q_{min}^* (e_{min}^*) is equal to 0.5 kg (500 TZS) for rice, 0.25 kg (250 TZS) for sugar, 1.05 kg (300 TZS) for tomatoes, and 0.35 kg (100 TZS) for onions. Note that in this district, the unit price schedule for rice is flat, so that losses on rice are zero by construction. The final column shows the percentage of transactions in the item-district group that are covered by the focal quantities. Coverage rates range from 37% to 91%; all but two are 69% or greater. Across all items and districts, three quarters of focal point coverage rates are above 50%.

Because price schedules are estimated at the district level, and the data are from seven districts, there are up to seven possible values of q_{min}^* and e_{min}^* for each item. Table VI shows the mean, minimum, and maximum values of these two statistics, across the seven districts. The table also includes the mean quantity purchased over 2 weeks, in column 1, for sake of comparison with q_{min}^* . In columns 3, 4, 6, and 7 we see substantial spatial variation in the minimum and maximum of q_{min}^* and e_{min}^* ; what it means to “buy in bulk” varies across districts. Nonetheless, among households that purchase each item, the mean quantity purchased (column 1) substantially exceeds the average value of q_{min}^* (column 3).⁶ This highlights an important point that we will discuss further in Section 5.1.2: in these data, the quantities required to buy in bulk are not especially large.

To provide a visual example from the data, Figure II depicts expenditures and unit prices by quantity for the 686 purchases of kerosene in one of the study districts. The size of the circles corresponds to the number of transactions at the circle center. The triangles represent the estimated focal points, and the solid lines mark the unit price (left panel)

⁶For many items, households would likely have access to even lower bulk discount prices if they were to buy large, wholesale quantities. Such quantities are not counted as focal by our method, because they are not represented by 5% or more of all purchase quantities. In this sense we are likely to underestimate the degree of available bulk discounts.

Table VI: Summary statistics across districts for q_{min}^* and e_{min}^* , by item

Item	Average quantity purchased	Across the 7 study districts...					
		q_{min}^*			e_{min}^*		
		Mean	Min	Max	Mean	Min	Max
(1)	(2)	(3)	(4)	(5)	(6)	(7)	
Maize	20.9	9.38	3	20	2414	750	6498
Kerosene	.95	.86	.5	1	1071	600	1300
Cooking Bananas	17.32	10.8	1.72	28	1067	201	1708
Cooking Oil	.84	.51	.05	1	850	100	1800
Rice	5.86	1.14	.5	4	771	300	2400
Sugar	2.22	.5	.25	2	629	250	2600
Flour	7.19	1.39	.25	4	617	100	1600
Charcoal	13.05	2.61	1.45	7.25	400	200	700
Beans	2.19	.43	.25	1	400	200	900
Coconut	3.88	.88	.57	1.1	383	200	550
Milling	20.57	9.17	4	20	349	100	700
Cassava	7.52	2.68	.58	8.67	325	50	997
Sweet Bananas	.97	1.57	.05	8.61	276	50	550
Salt	1.06	.57	.25	1	264	100	500
Tea Leaves	.08	.08	.01	.25	243	50	500
Soap	6.56	2.29	1	8	243	100	704
Sardines	1.22	.27	.14	.5	186	100	300
Tomatoes	2.85	.65	.35	1.4	171	100	300
Matches	4.18	3.57	1	10	133	30	400
Onions	1.15	.51	.05	1.4	114	50	200
Tea	3.32	.5	.5	.5	100	100	100
Cigarettes	30.69	2	1	6	91	50	240

Notes: Authors' calculations from SHWALITA data. Column 1 refers to average total purchase over 2 weeks at the household-item level, for households that purchased positive amounts of the item. Table is sorted by column 5. Units listed in Table I.

and expenditure (right panel) schedules. In the left side panel, the downward orientation of the unit prices is clear. In the right side panel, the changing slope of the expenditure line represents the drop in unit prices as quantity increases. The clustering of purchases at a small number of quantities is also clear.

We have shown a general pattern of bulk discounts in the data. However, it is important to note that in each district there are also some items that do not exhibit bulk discounts. Estimated unit price schedules are flat for 64 of the 146 item-district groups in the




Figure II: Expenditure and unit price for kerosene purchases in one district

Notes: Authors' calculations from SHWALITA data. We dropped 11 outliers to improve readability of these figures.

data. Some items exhibit bulk discounts in every district. These include maize, cooking oil, kerosene, cooking bananas, and tea leaves. Unit price schedules for other items, including sweet bananas, cooking bananas, onions, salt, milling, and sardines, are downward-sloping in a majority of districts. In contrast, brewed tea never exhibits bulk discounts, while cigarettes, beans, rice, cassava, and matches each only exhibit discounts in 1 or 2 districts. In Section 5.2.1 we will exploit the fact that households purchase from both flat and downward-sloping unit price schedules to study the salience of discounts at the time of purchase.⁷

4.2 The value of forgone consumption

We turn now to our key welfare measures: the quantity and value of forgone consumption from buying in small quantities. Recall from section 2.2 that the quantity of consumption forgone is given by $Q_{hi} = q_{hi}^* - q_{hi}$, where the first term is the inverse of the expenditure function evaluated at total adjusted expenditure e_{hi}^* , and the second term is total observed quantity. Likewise, the financial loss, or value of forgone consumption, is defined as $L_{hi} =$

⁷Why discounts emerge for some items and not for others, and why they persist despite the appearance of robust competition in retail markets, are open questions not addressed in this paper. See Attanasio and Pastorino (2015).

$\hat{e}_{hi} - e_{hi}^*$, where \hat{e}_{hi} is total adjusted expenditure and e_{hi}^* is the cost of buying q_{hi} in a single transaction. Summing across items at the household level gives $L_h = \sum_i L_{hi}$.

We begin with the quantity measures. In columns 1–3 of Table VII we report the item-level means of total observed quantity, q_{hi} , counterfactual quantity, q_{hi}^* , and the counterfactual increase in quantity in percentage terms, \tilde{Q}_{hi} . Calculations in this table are based on all households that purchase an item more than once. The results are striking: without changing total expenditure, households could increase quantity purchased by almost 16% on average. Potential quantity increases are over 25% for kerosene, onions, cooking bananas, cooking oil, tea leaves, and sardines. These are staple goods: kerosene is the primary lighting fuel in much of Tanzania, cooking bananas are a staple carbohydrate (in the two districts where they are most commonly purchased), sardines are a key source of protein, and cooking oil is the main source of cooking fat. Most households purchase one or more of these goods: 85% purchase kerosene, 78% purchase cooking oil, 70% purchase sardines, and 21% purchase cooking bananas (Table II). By choosing to spend \hat{e}_{hi} in small increments, the average household is sharply reducing its consumption of these staple items. The welfare losses implied by columns 1–3 of Table VII, then, are substantial at face value.

The money-metric measures of loss tell a similar story. In columns 4–6 of Table VII we report summary statistics for \hat{e}_{hi} , L_{hi} , and \tilde{L}_{hi} . (The table is sorted by decreasing values of column 6, so that high loss items are at the top. From now on we will usually display items in that order.) On average, losses represent 8.9% of total expenditure at the household-item level. For a number of frequently purchased items – sardines, onions, kerosene, cooking oil, cooking bananas, tomatoes – losses represent more than 10% of expenditure. In columns 4–6 of the lower panel of Table VII we report summary statistics for all households represented in the upper part of the table, divided into those above/below median L_h . We calculate household-level means by first summing adjusted expenditure ($\hat{e}_h = \sum_i \hat{e}_{hi}$) at the household level, then averaging. We define the household-level percentage loss measure, \tilde{L}_h , as $\tilde{L}_h = L_h/\hat{e}_h$. Not surprisingly, losses vary substantially across households. The overall household-

Table VII: Purchase quantities and expenditures: Observed and counterfactual

	(1)	(2)	(3)	(4)	(5)	(6)
ITEM-LEVEL MEANS						
Item	Quantity q_{hi}	Potential quantity q_{hi}^*	% change \tilde{Q}_{hi}	Adjusted expenditure \hat{e}_{hi}	Loss L_{hi}	% Loss \tilde{L}_{hi}
Kerosene	1.12	1.33	33.3	1730	289	19.8
Onions	1.42	1.81	46.1	522	95	19.7
Cooking Bananas	26.65	31.06	50.7	3044	496	18.4
Cooking Oil	0.99	1.19	26.4	2149	311	16.5
Tea Leaves	0.09	0.11	31.5	482	57	16.1
Sardines	1.53	1.71	28.6	903	113	13.8
Tomatoes	3.39	3.89	17.5	1244	107	11.1
Salt	1.41	1.48	7.5	671	37	6.9
Coconut	4.53	4.80	7.7	1889	115	6.8
Maize	34.12	36.12	8.3	11846	711	6.6
Sweet Bananas	1.47	1.55	6.3	1157	31	4.7
Cassava	11.73	12.14	5.6	1355	66	4.6
Soap	8.58	8.88	4.5	925	36	4.4
Charcoal	14.88	15.37	7.2	5465	56	4.2
Milling	29.89	30.56	3.0	856	22	2.8
Matches	5.44	5.56	2.5	248	5	2.2
Cigarettes	40.40	40.55	0.4	1955	14	1.2
Sugar	2.70	2.71	0.8	3212	19	1.1
Flour	8.76	8.82	0.7	4944	38	1.0
Rice	7.55	7.59	0.6	6385	43	0.9
Beans	2.81	2.81	0.5	2809	9	0.5
Tea	4.53	4.53	0.0	905	0	0.0
AVERAGE			15.6	2106	112	8.9
HOUSEHOLD-LEVEL MEANS						
Households				\hat{e}_h	L_h	\tilde{L}_h
All				15826	840	6.9
Below median				9290	198	4.5
Above median				22335	1480	9.3

Notes: All multi-purchasing households included; L_{hi} and percent change both set to zero for single-purchasing households; for Item panel, columns 3 and 6 calculated at household-item level before averaging across items, and column 3 calculated after throwing out upper 1% tail; for Household panel, “median” refers to median of L_h .

level average is 840 TZS, or 6.9% of expenditure. Financial losses among the above median group represent almost 10% of total expenditure, on average.

Figure 4.2 shows histograms and kernel density estimates for the distributions of L_h (left panel) and \tilde{L}_h (right panel) among multi-purchase household-item pairs. Items with flat price schedules are not dropped, so as not to bias the estimates toward large losses. There is substantial between-household variation in losses. Approximately 8% of households incur zero losses (with our conservative approach to estimation). Yet, nearly a quarter (24%) incur




Figure III: Distribution of financial losses

Notes: Authors' calculations from SHWALITA data. The average exchange rate during the study period was 1,150 TZS per US dollar.

losses above 10% of expenditure.

What drives the substantial between-household variation in financial losses? To examine whether particular types of households are more prone to forego bulk discounts, we estimate household-level descriptive regressions of L_h and \tilde{L}_h on a vector of household characteristics. Results are shown in Table VIII. Columns 1 and 2 report the estimates with L_h as the dependent variable. In both columns we see that the poorest quartile of households (the excluded category) have lower losses than the other three quarters of households, though the differences are only weakly statistically significant. The age, gender, and education level of the household head do not meaningfully co-move with losses. Larger households exhibit slightly greater losses, a result we discuss further in Section 5.2.4.

Columns 3 and 4 of Table VIII report the results of the same specifications, with \tilde{L}_h as the dependent variable. Households in the first two wealth quartiles have similar mean percentage losses, but percentage losses are slightly lower for those in the third quartile, and substantially lower for those in the wealthiest quartile. The estimated coefficients on head of household characteristics, in column 4, are too small in magnitude to be of importance. While distance from the community center is statistically significantly associated with lower

Table VIII: Loss regressed on household characteristics

Dependent variable:	Loss	Loss	% Loss	% Loss
	(1)	(2)	(3)	(4)
Wealth index quartile 2 (=1)	162.752*	153.190*	-0.002	-0.001
	(78.307)	(72.942)	(0.006)	(0.006)
Wealth index quartile 3 (=1)	153.415	140.203	-0.010**	-0.010**
	(107.081)	(126.721)	(0.004)	(0.003)
Wealth index quartile 4 (=1)	270.419*	228.306	-0.033**	-0.034***
	(116.184)	(158.872)	(0.010)	(0.009)
Age of head (years)		-5.320*		-0.001**
		(2.352)		(0.000)
Head is female (=1)		29.313		0.003
		(78.002)		(0.005)
Head years of education		-0.908		-0.001
		(15.013)		(0.001)
Household size		78.799*		-0.004
		(32.241)		(0.003)
Distance to community center (km)		-91.024		-0.006**
		(70.775)		(0.002)
Observations	1471	1465	1471	1465
R-squared	0.16	0.17	0.19	0.21
Mean dep. var.	840	837	0.069	0.069

Notes: Authors' calculations from SHWALITA data. Standard errors in parentheses; standard errors clustered at district level; *** sig. at 0.01, ** sig. at 0.05, * sig. at 0.1. All regressions include district fixed effects, controls for demographic composition of the household, and controls for questionnaire module. The wealth index is defined with quartile 1, the excluded group, as the poorest.

normalized losses, the effect is not especially meaningful given the distribution of distances in the data (mean 0.6 kilometers, s.d. 0.7).

Perhaps the most interesting results in Table VIII are those related to wealth. When using levels (columns 1-2), losses appear to be positively related to wealth. When using percentages, losses are negatively related to wealth. This pattern indicates that there may be different types of loss-prone households – wealthy households that suffer large losses in levels but small losses as a percentage of total spending, and poor households that suffer small losses in levels but large losses as a percentage of total spending. To investigate this possibility, Table IX presents summary statistics for four groups of households (moving from column 2 to column 5): (i) households in the highest quartile of L_h but not the highest quartile of \tilde{L}_h , (ii) households in the highest quartile of \tilde{L}_h but not the highest quartile of L_h , (iii) households in the highest quartile for both losses and percentage losses, and (iv)

households that are in neither worst quartile.

Table IX: Summary statistics by loss categories, household level

	Among the 25% highest loss households by...				
	Overall	L_h only	\tilde{L}_h only	Both L_h and \tilde{L}_h	Neither
	(1)	(2)	(3)	(4)	(5)
Proportion in group	1.00	0.12	0.12	0.13	0.63
Number of transactions	37.95	79.53	23.95	49.32	30.49
Number of items purchased	10.65	14.36	8.93	12.15	9.96
Adjusted expenditure	17589	40975	6186	19299	14985
Adjusted expenditure per capita	4060	8447	2055	4311	3565
Loss (level)	825.54	1944.99	658.04	2283.98	343.85
Loss (%)	0.05	0.05	0.11	0.13	0.03
Wealth index	-0.01	0.83	-0.41	-0.01	-0.10
Distance to comm. center (km)	0.61	0.36	0.57	0.53	0.67
Head age (years)	46.66	46.44	46.25	41.13	47.98
Head education (years)	4.73	5.77	3.78	5.17	4.60
Head is female (=1)	0.20	0.19	0.26	0.17	0.20
Household size	5.33	6.14	4.32	5.25	5.39
Number of children 9-14	0.91	1.14	0.65	0.87	0.92
Number of adults 15-59	2.50	3.06	1.99	2.43	2.50

Notes: Authors' calculations from SHWALITA data. Sample includes the 1,497 households with complete demographic data.

Groups (i) and (ii) look like the rich and poor households discussed in the previous paragraph. The 12% of households that have high losses but not high percentage losses (column 2) appear to be upper-class households. They make substantially more purchases, spend more than twice as much, and buy many more items than the average household. Their average level of the wealth index is almost a full standard deviation above the mean, and they are larger, more educated, and live nearer to the city center. In contrast, the 12% of households that have high percentage losses but not high level losses (column 3) appear to be poor and disadvantaged households. These households are smaller, less educated, and poorer in both expenditure and wealth terms.

Group (iii), the 13% of households that are in both high-loss categories (column 4), are interesting for a different reason: they exhibit very large losses despite having close to average expenditures. They also have near average wealth, household size, and education. Their most notable characteristic is that the household heads are younger and more likely to

be male, raising the interesting possibility that they lack the foresight or maturity to organize household finances. Otherwise, there is little besides their inefficient shopping patterns that distinguishes these households from the average.

When we investigate in Section 5 why some households forego bulk discounts, we will often use the categorization of households from Table IX. Indeed, given how different these three groups of loss-prone households appear to be, it seems plausible that they suffer losses for different reasons.

4.3 Variation in price paid conditional on quantity

Although our focus is on bulk discounts and why households might not take advantage of them, it is worth taking a moment to explore the nature of residual variation around the expenditure schedule. As suggested in the motivating example in Section 2.2, we observe many instances in which the price for the same quantity of the same item varies between transactions. Because of how we construct L_h and \tilde{L}_h , this price variation does not impact our loss analysis directly. But it is interesting, nonetheless, because it represents a second dimension of between-household variation in prices that may be responsible for the “poor pay more” hypothesis.

In Table X we show the proportion of transactions for each item that are below, on, and above the expenditure schedule. There is less variation than one might expect. On average, 46% of transactions are exactly on the schedule, with 19% below and 35% above. At the top of the table, with 74–95% of prices falling on the schedule, we find matches, tea and cigarettes. These are highly standardized goods that are sold in clearly identifiable and uniform units. At the bottom of the list are cooking bananas and cassava, with less than 20% of transactions on the schedule. These goods are typically sold in imprecise units (heaps, bunches). This suggests that some of the variation in unit price conditional on quantity may be due to measurement error, either at the time of purchase or during data collection.

In Section 2.2 we labeled the idiosyncratic component of price, conditional on quan-

Table X: Position of transaction expenditure relative to expenditure schedule

Item	Below (1)	On (2)	Above (3)
Cigarettes	0.01	0.95	0.04
Tea	0.10	0.78	0.12
Matches	0.14	0.74	0.12
Sugar	0.17	0.70	0.13
Onions	0.17	0.61	0.21
Soap	0.11	0.57	0.32
Tomatoes	0.18	0.52	0.30
Rice	0.20	0.49	0.31
Tea Leaves	0.11	0.48	0.41
Beans	0.28	0.44	0.28
Salt	0.18	0.42	0.40
Kerosene	0.26	0.39	0.36
Charcoal	0.26	0.38	0.36
Dagaa	0.13	0.35	0.52
Cooking Oil	0.20	0.35	0.46
Coconut	0.33	0.32	0.34
Sweet Bananas	0.17	0.30	0.53
Maize	0.27	0.25	0.48
Flour	0.21	0.22	0.57
Milling	0.20	0.20	0.61
Cooking Bananas	0.32	0.17	0.50
Cassava	0.47	0.12	0.41
AVERAGE	0.19	0.46	0.35
Wealth index quartile 1	0.17	0.48	0.35
Wealth index quartile 2	0.18	0.46	0.35
Wealth index quartile 3	0.20	0.47	0.33
Wealth index quartile 4	0.20	0.45	0.35

Notes: Authors' calculations from SHWALITA data. The wealth index is defined with quartile 1 as the poorest. Table sorted by decreasing values of column 2.

tity, as ν_{hik} . This residual variation could reflect unobserved item quality, bargaining skill, shopping effort, or other factors. We can calculate the empirical analog of this term as the difference between observed and adjusted expenditure, i.e., $\hat{\nu}_{hik} = e_{hik} - \hat{e}_{hik}$. By definition, the 46% of transactions that take place on the expenditure schedule have $\hat{\nu}_{hik} = 0$. To examine the correlates of $\hat{\nu}_{hik}$, we first normalize it to its percentage difference from the expenditure schedule: $\hat{\nu}_{hik}^n = \hat{\nu}_{hik} / \hat{e}_{hik} = \frac{e_{hik} - \hat{e}_{hik}}{\hat{e}_{hik}}$, where the “n” superscript indicates “normalized.” The mean of $\hat{\nu}_{hik}^n$ is 0.16, indicating that the average transaction is 0.16 standard deviations above the expenditure schedule.⁸

⁸Recall that the focal expenditures that underlie the expenditure schedule are medians. The average transaction lies above the schedule because there is positive skewness in expenditure conditional on quantity.

Table XI: Regressions with idiosyncratic price component as dep. variable, transaction level

Dependent variable:	$\hat{\nu}_{hik}^n$ (1)	$ \hat{\nu}_{hik}^n $ (2)
Quantity z-score	-0.029 (0.02)	-0.022 (0.02)
Precise unit (=1)	0.112 (0.10)	0.116 (0.10)
Market day purchase (=1)	0.013*** (0.00)	0.007 (0.01)
Wealth index quartile 2 (=1)	-0.017 (0.02)	-0.011 (0.02)
Wealth index quartile 3 (=1)	-0.002 (0.02)	0.001 (0.01)
Wealth index quartile 4 (=1)	0.032 (0.02)	0.021 (0.02)
Observations	53588	53588
R-squared	0.20	0.23
Mean of dep. variable	0.16	0.24

Notes: Authors' calculations from SHWALITA data. Standard errors in parentheses; standard errors clustered at district level; *** sig. at 0.01, ** sig. at 0.05, * sig. at 0.1. All regressions include district fixed effects, item fixed effects, and controls for questionnaire module. Sample includes 1,496 households in 168 villages in 7 districts. We dropped observations in the 1% upper and lower tails of the $\hat{\nu}_{hik}^n$ distribution before estimation.

To examine the variation in $\hat{\nu}_{hik}^n$, we estimate regressions of the level and absolute value of $\hat{\nu}_{hik}^n$ on transaction and household characteristics. We use both the level and absolute value as dependent variables so as to explore factors associated with higher prices and greater spread. Table XI shows results. The variables of main interest are the wealth quartile dummies and the variable “Precise unit”, which takes a value of 1 if the unit involved in the transaction is standardized and precisely defined (at the local level), and zero otherwise.⁹ Regressions also include district effects, item effects, questionnaire effects, controls for quantity (via item-level z-scores), and controls for purchases on village market days.

Results are broadly similar across the two columns of Table XI. There is only one statistically significant coefficient, in column 1, indicating that average prices are conditionally

⁹Based on the market survey efforts of the research team, we designated the following units as precise: kilogram, liter, 25kg bag, 50kg bag, *debe*, *kisadolini*, and packet of tea leaves. These units are associated with standardized quantities that were measured by the research team at markets in every village. Imprecise units include bowls, cups, pieces, heaps, and others. These were also surveyed and measured by the research team, but they are prone to greater measurement error. Approximately 63% of transactions were recorded in precise units.

higher on market days by a small amount. The “Precise unit” variable is not statistically different from zero, and has the opposite sign of that expected.¹⁰ Otherwise, the main take-away is that the residual component of prices does not vary meaningfully with wealth. The estimated coefficients on the wealth quartile dummy variables are neither economically nor statistically significant. This establishes the main result for this subsection: on average, there do not appear to be unobserved transaction characteristics that lead to poor households paying different prices from wealthy households for the same quantity of the same item.

5 What explains the observed purchasing patterns?

We now turn to the mechanisms that might lead households to engage in financially inefficient purchasing patterns by foregoing bulk discounts. We consider a number of possibilities, as previewed in Section 2. We collect these mechanisms into two groups: those that we argue we can reject (Section 5.1), and those that we think might be playing a significant role (Section 5.2).

5.1 Mechanisms that we reject

Our analysis leads us to reject the following mechanisms as reasons for the observed purchasing patterns: binding liquidity constraints, storage and transport costs, ignorance of bulk discounts, utility from frequent shopping, and maintenance of buyer-seller relationships. By “reject” we do not mean that these mechanisms are irrelevant for all households. Rather, we do not think these mechanisms are responsible to any substantial degree for the losses that we identify in Section 4.

¹⁰If we exclude item fixed effects from these regressions, the “Precise unit” coefficient is negative, larger in magnitude, and borderline statistically significant. The implication is that if measurement error matters, the effect is not distinguishable from between-item variation in price conditional on quantity.

5.1.1 Liquidity constraints

Perhaps the most natural explanation for the losses in our data is that people would like to take advantage of bulk discounts, but they lack the liquidity to do so. Prior work has emphasized that poor households may pay higher unit prices than wealthy households because binding liquidity constraints prevent them from taking advantage of bulk discounts.

Initial evidence against this hypothesis is in Table VIII, above. There, in column 1 we see that wealthier households, who would seem least likely to be subject to liquidity constraints, do indeed incur losses. To dig deeper, we investigate whether poorer households generally buy in smaller quantities than wealthier households, which would be consistent with the liquidity constraints hypothesis. Table XII shows the average purchase quantities for each item, by wealth quartile. Wealth quartiles are defined within districts, to match the expenditure schedules. Items are ordered by decreasing values of average percentage losses (from column 6, Table VII). Looking across rows of the table, it is clear that the average quantities in the first wealth quartile are not generally smaller than those in the fourth quartile. For two of the first four goods listed, the poorest quarter of households buy the largest quantities, on average. Across the remaining goods there is no clear ranking between the wealth quartiles. In combination with the finding in Table XI that the poor do not pay higher prices than the wealthy for the same quantities, there is no evidence of the poor paying higher prices because they cannot take advantage of bulk discounts.

The rich-vs.-poor analysis above assumes that poorer households are more likely to be liquidity constrained. Yet it is possible that many households in the data, including many wealthy households, are liquidity constrained. To approach the question of liquidity constraints from a different angle, we ask the following: for how many days would a household have to delay purchasing an item in order to buy it at the lowest available unit price (which it could then do in perpetuity)? In other words, for how long would a household need to forego consumption in order to overcome a liquidity constraint? Let $a_{hi} \equiv \hat{e}_{hi}/14$ be the average daily expenditure on item i by household h , and recall that e_{min}^* is the minimum expenditure

Table XII: Average quantity per transaction, by item

Item	Wealth Quartile (1 = poorest)			
	1	2	3	4
Kerosene	0.22	0.23	0.27	0.35
Onions	0.33	0.30	0.28	0.31
Cooking Bananas	10.95	7.59	6.99	5.82
Cooking Oil	0.16	0.16	0.18	0.24
Tea Leaves	0.02	0.02	0.02	0.03
Sardines	0.36	0.32	0.38	0.39
Tomatoes	0.56	0.56	0.56	0.60
Salt	0.51	0.52	0.50	0.55
Coconut	0.89	0.76	0.76	0.67
Maize	10.22	11.27	9.02	11.23
Sweet Bananas	0.40	0.47	0.35	0.45
Cassava	3.50	2.79	3.15	3.23
Soap	2.05	2.22	2.18	2.41
Charcoal	2.09	2.01	2.02	2.10
Milling	8.57	8.29	7.71	10.03
Matches	1.98	1.78	1.90	2.26
Cigarettes	4.83	4.34	5.90	5.34
Sugar	0.51	0.48	0.51	0.63
Flour	1.38	1.28	1.18	1.27
Rice	1.68	1.57	1.61	1.64
Beans	0.88	0.93	0.82	0.79
Tea	0.67	0.77	0.74	0.78

Notes: Authors' calculations from SHWALITA data. See Table I for units. Items ordered by decreasing values of mean \tilde{L}_h , from column 6 in Table VII.

required to buy item i at the lowest focal unit price. The self-financed purchasing delay is then $d_{hi} = e_{min}^*/a_{hi}$. We calculate d_{hi} for all household-item pairs in which the household makes at least one transaction.¹¹ This is a very conservative approach, because it ignores the fact that households could also shift spending between items, and might be able to access some credit to finance the initial bulk purchase.

In Table XIII, we report the item-level median value of d_{hi} for all households (column 1), by wealth quartile (columns 2-5), and for the three groups of loss-prone households (columns 6-8). The most striking feature of Table XIII is that the median delays are all relatively short. Virtually all are less than two weeks, and the majority are less than one week. The wealthiest households have shorter delays, but the differences are not substantial.

¹¹Note that including household-item pairs with only one transaction will tend to make these numbers larger because this will include someone who is just dabbling—i.e., someone who makes one small quantity purchase will have a high value of d_{hi} .

Table XIII: Median days required to save enough to purchase at lowest unit price

Item	All	By wealth quartile (1 = poorest)				Loss-prone HHs as measured by...		
	(1)	1	2	3	4	%age only	Level only	Both
Kerosene	14.0	17.8	14.0	13.7	10.1	19.9	7.8	13.4
Onions	4.7	5.6	4.7	4.7	3.5	7.0	3.2	3.9
Cooking Bananas	11.9	14.0	12.9	14.0	9.3	24.1	9.3	8.0
Cooking Oil	8.1	13.2	11.5	8.9	3.7	15.0	4.2	7.2
Tea Leaves	7.9	7.0	14.0	14.0	4.7	35.0	4.5	16.3
Sardines	4.7	3.5	3.9	4.7	7.0	5.2	4.1	4.0
Tomatoes	3.1	4.7	4.7	3.5	1.7	4.7	1.7	2.5
Salt	8.6	7.0	7.0	8.8	11.7	11.7	7.8	8.8
Coconut	3.3	2.8	3.1	2.8	3.7	11.2	3.1	5.6
Maize	7.2	7.0	9.8	8.0	12.3	17.5	4.7	8.7
Sweet Bananas	8.4	10.5	16.8	9.5	6.0	18.2	4.8	13.2
Cassava	5.3	4.7	3.1	3.5	9.8	6.1	8.1	5.7
Soap	4.7	6.2	4.0	3.5	7.0	7.0	3.5	3.5
Charcoal	2.0	14.0	7.0	3.5	1.7	12.9	1.6	3.5
Milling	5.8	5.6	4.7	7.1	12.3	7.8	4.1	7.0
Matches	7.0	7.0	7.0	12.4	7.0	14.0	7.0	7.0
Cigarettes	1.0	1.6	1.3	0.9	0.5	3.5	0.5	1.3
Sugar	2.3	3.5	2.8	2.3	1.6	4.5	1.4	2.8
Flour	1.6	1.0	1.2	1.3	2.3	3.5	1.4	2.3
Rice	2.3	3.5	3.5	3.1	1.0	7.0	1.1	4.7
Beans	3.5	3.5	3.5	3.5	1.8	4.7	1.8	3.5
Tea	3.5	7.0	3.5	4.7	2.0	7.0	2.3	3.5
OVERALL	4.7	5.8	6.0	5.3	3.5	8.8	3.4	5.4

Notes: Authors' calculations from SHWALITA data. Items ordered by decreasing values of mean \tilde{L}_h , from Table VII. Column 6 includes households in highest quartile by \tilde{L}_h but not L_h . Column 7 includes households in highest quartile by L_h but not \tilde{L}_h . Column 8 includes households in highest quartile by both.

The “Percentage-only” loss-prone households have the longest delays, with an average of 8.8 days (column 6). We have already seen in Table IX that these households appear to be poor and disadvantaged on numerous dimensions. Hence, if liquidity constraints drive small-quantity purchasing for anyone, these households are the most likely candidates. Yet, even for this group, the delays are short. If liquidity constraints are the *only* barrier to high-quantity purchasing, then it hardly seems onerous to delay consuming an item for a week or two, once, in order to consume more of it at lower cost forever afterwards.

Mullainathan and Shafir (2013) conduct a similar analysis for roadside vendors in Chennai, India (p. 123-124). Those vendors lose roughly half of their daily earnings to inter-

est payments on short-term loans, and yet still buy a daily cup of tea. The authors calculate that by foregoing tea for 50 days, the average vendor could save enough to permanently avoid short-term borrowing, resulting in doubling of take-home pay (and, from then on, the purchase of more tea in perpetuity). Here, we find far shorter delays are required to reach the efficient purchasing path, even for the households with the highest losses. Furthermore, the analysis in Table XIII is purely within-good. If we were to allow for subsidization across goods, more in line with the intuition in Mullainathan and Shafir (2013), the delays would be even shorter.

From this evidence, we conclude that liquidity constraints are not playing a major role. This is not to say we believe liquidity constraints are irrelevant for all households. Rather, the losses that we observe in our data do not seem to be driven by liquidity constraints.

5.1.2 Costs of transporting or storing bulk purchases

In wealthy countries, limitations on transport or storage space can be a binding constraint on bulk purchasing (Griffith et al., 2009). A household cannot buy a carton of paper towels at a big box store if it has no way to transport and store such a large purchase. In our data, the situation is not so extreme. The bulk discounts in our study are available at relatively small quantities, almost all of which can easily be transported and stored in a typical household.

To see this, compare in Table VI the distribution of q_{min}^* (columns 2-4) with the mean total purchase quantities across two weeks (column 1). The magnitudes in column 2 are generally far below those in column 1. Total household purchase quantities over just two weeks typically exceed the minimum quantity needed in order to buy in bulk. In this setting, bulk purchasing does not require the purchase and transport of an excessive quantity of any item.

The same line of reasoning makes it very unlikely that households avoid bulk purchasing because of concerns about theft or depreciation of stocks during storage. In Table VI we see that the average household buys many multiples of the minimum quantity needed

to access the lowest available unit price. The quantities associated with bulk purchasing are simply too small for depreciation to be the central concern of consumers.

Finally, if transport costs were a major impediment to bulk purchasing, we would expect households living on the outskirts of the community to have higher losses. Yet, in columns 2 and 4 of Table VIII we see that the association between distance and loss is negative, if significant at all. Overall, there is no evidence that the cost of transporting or storing bulk purchases can explain the purchase quantities in the data.

5.1.3 Utility from shopping

Perhaps people make frequent, small-quantity purchases because there is a utility value from shopping—e.g., from the socializing and community engagement that one enjoys in the market. Given our knowledge of Tanzania, we doubt this mechanism plays much of a role in leading people not to take advantage of bulk discounts. People can pass by shops and markets without making purchases, and do so frequently. There is nothing to stop consumers from visiting shops to socialize, but buying in bulk when they need to make a purchase.

The data also contradict this hypothesis. Even the households that incur the largest losses could rearrange their purchase patterns to take greater advantage of bulk discounts *and* make more purchases overall. To illustrate, define $K_{hi}^* \equiv \hat{e}_{hi}/e_{min}^*$. This is the counterfactual number of separate transactions that household h could make on item i at the lowest available unit price for that item. The actual number of transactions that household h makes on item i , K_{hi} , could be smaller or larger than K_{hi}^* . For items with bulk discounts on which the households makes frequent small purchases—in particular, purchases of quantities smaller than q_{min}^* —the household will have $K_{hi}^* < K_{hi}$. For items on which the household typically purchases quantities larger than q_{min}^* , the household will have $K_{hi}^* > K_{hi}$. For items purchased in both small and large quantities, the ordering of K_{hi}^* and K_{hi} is ambiguous.

Table XIV shows the mean values of K_{hi} , K_{hi}^* , and the difference $K_{hi}^* - K_{hi}$ for all households and for the three groups of loss-prone households. The positive differences in

Table XIV: Counterfactual change in number of transactions, if purchasing at lowest unit price

Subgroup	Statistic	(1)	(2)	(3)
		Actual (K_{hi})	Counter- factual (K_{hi}^*)	Difference
All households	Mean total transactions	38.0	61.3	23.3
	Mean transactions per item	3.6	5.8	2.2
Loss-prone households, level only	Mean total transactions	76.8	113.0	36.2
	Mean transactions per item	5.4	8.0	2.6
Loss-prone households, %age only	Mean total transactions	21.0	22.8	1.8
	Mean transactions per item	2.4	2.6	0.2
Loss-prone households, both	Mean total transactions	47.1	51.3	4.2
	Mean transactions per item	4.0	4.3	0.4

Notes: Authors' calculations from SHWALITA data. Second group of households are in highest quartile by L_h but not \tilde{L}_h . Third group are in highest quartile by \tilde{L}_h but not L_h . Fourth group are highest quartile by both.

column 3 indicate that the average household in all four groups could shop more while only buying in bulk. The transaction counts are consistent with what we already know about the loss subgroups (from Table IX). Households that are loss-prone in levels (only) accrue their high losses by shopping a lot, making 76.8 transactions on average. However, these loss-prone households could shop substantially more while paying the lowest unit price – 36.2 more transactions on average – if they eliminated transactions both below and above q_{min}^* . The next group, the high percentage loss group, are relatively poor and tend to shop very little, which is reflected in their low level of average transactions (21). Yet, even this group could make 0.2 more transactions per item, on average, while only purchasing in bulk. Finally, households that are in the 25% worst group in both levels and percentages could make 51.3 instead of 47.1 transactions if their goal was to shop as much as possible while never paying more than the lowest available unit price.

Table XIV does not show the between-item variation in K_{hij}^* and K_{hi} .¹² There are some high-loss items that would have to be purchased less often by all groups (kerosene stands out). On net, however, the potential increases in transactions for most items outweigh the required cutbacks in others.

¹²See online appendix Table A.2 for the item-level breakdown underlying Table XIV.

These counterfactual shopping patterns are unlikely to be optimal for a variety of reasons. However, what this analysis clearly demonstrates is that the desire to shop frequently cannot explain the failure to take advantage of bulk discounts, because households could already do more of both.

5.1.4 Ignorance of bulk discounts

Is it possible that many people in Tanzania simply do not know of the available bulk discounts? We are doubtful. When we conducted informal interviews with individuals in the study area, everyone was well aware of bulk discounts for a wide range of consumption items. Furthermore, our data suggest that many households personally experience the non-linear prices. Column 2 of Table IV shows the results of item-level regressions of unit price on quantity, with household fixed effects. Even *within household*, bulk discounts show up. The members of households that purchase items multiple times – exactly those that are foregoing potential consumption – are surely aware of the available discounts.

5.1.5 Investment in buyer-seller relationship

Could it be that bulk discounts are only available to consumers who have an established relationship with a vendor? In this case, a buyer might pay higher unit prices today as an investment in a relationship that will allow future access to better prices. Perhaps our data reflect a point-in-time snapshot of an ongoing process in which consumers gradually cultivate, maintain, and sometimes lose these vendor relationships? Or, relatedly, it may be that vendors are only willing to sell some items as “loss leaders” – large quantity purchases provided at a heavy discount – when they are combined with smaller quantity purchases at higher unit prices. If either of these scenarios holds, then our measures of financial loss would be incorrect, because the expenditure schedules that we estimate from the data would not reflect a genuine set of counterfactuals.

However, the data collected by project staff members clearly contradicts this hypoth-

esis. In their market price surveys, bulk discounts are clearly present (see columns 4 and 5 of Table IV). Yet, these staff members had no prior relationship with vendors. Clearly, consumers do not need to invest in long-term relationships with sellers, nor must they combine large and small quantity purchases, in order to receive bulk discounts.

5.2 Mechanisms consistent with the data

We next turn to mechanisms for which we find supporting evidence. These are: inattention to the magnitude of forgone consumption, purchasing in small quantities as a form of self- or other-control, avoidance of social taxation, and coordination costs within the household.

5.2.1 Inattention

In Section 5.1.4 we argued that ignorance of the available bulk discounts does not drive our findings. But it consumers could still be inattentive to the magnitude of the financial benefits (or extra consumption) from taking greater advantage of bulk discounts, and, as a result, not react to the existence of bulk discounts. To investigate this possibility, we exploit the fact that households face some flat and some non-flat unit price schedules. In particular, for the average study household, roughly 56% of observed purchases are of items with bulk discounts in local markets, and 44% are of items with flat price schedules. Using this variation, we can assess the extent to which households react to the existence of bulk discounts.

In order to implement this analysis across items, we use as the dependent variable K_{hi} , the number of transactions by household h for item i . If a household is attending to the impact of bulk discounts, then, relative to purchase patterns for goods with flat price schedules, the household should make fewer transactions—that is, purchase larger quantities less frequently. In other words, we take behavior on flat price schedules as reflective of baseline purchase patterns, and then assess the extent to which households adjust those patterns in reaction to the financial incentives created by bulk discounts. In this analysis it

is also important to control for other unobserved differences in the purchase patterns across households and items. Because there is within-item and within-household variation in the presence of bulk discounts, we can identify the key interaction terms while including both household and item fixed effects.

Columns 1 and 2 of Table XV show the results of regressing K_{hi} on a dummy variable for whether the item exhibits bulk discounts (columns 3 and 4 are discussed in the next subsection). In column 1, we see that households make fewer purchases on items with bulk discounts. Hence, on average, households indeed seem to attend to the existence of bulk discounts. However, column 2 investigates how this effect differs for the three categories of loss-prone households described in Section 4.2. The message is somewhat different. The non-loss-prone households exhibit a strong negative relationship between the number of transactions and the existence of bulk discounts. In contrast, loss-prone households in group (i) (in levels only) and group (ii) (in percentages only) exhibit no relationship between the number of transactions and the existence of bulk discounts, as indicated by the F-tests listed in the lower panel of the table. In other words, these households seem not to attend to the existence of bulk discounts.

Finally, loss-prone households in group (iii) (in both levels and percentages) actually make 0.947 *more* transactions on items with bulk discounts (seen from summing coefficients (a) and (d) in column 2). This result is surprising, although it does explain this group's particularly large losses. One possible cause of this purchasing pattern would be an item-level correlation between bulk discounts and some other characteristic that generates frequent purchases. We discuss one possible characteristic, how tempting it is to over-consume a good, in the next section.

Overall, we conclude from the analysis in column 2 that inattention to the financial implications of bulk discounts is an important driver of losses in our data. The low-loss households react to bulk discounts; the high-loss households do not. The behavior of the loss-prone households is consistent with other recent work on the lack of salience of total prices to

Table XV: Number of transactions and bulk discounts, household-item level

Dependent variable: number of transactions at household-item level				
	(1)	(2)	(3)	(4)
(a) Bulk discounts (=1)	-0.275*** (0.07)	-0.666*** (0.09)		-0.662*** (0.09)
(b) Bulk discounts \times Highest quarter levels only (=1)		0.455*** (0.15)		0.717*** (0.20)
(c) Bulk discounts \times Highest quarter %age only (=1)		0.827*** (0.15)		0.602*** (0.18)
(d) Bulk discounts \times Highest quarter, both (=1)		1.619*** (0.19)		1.361*** (0.23)
(e) Temptation good (=1)			0.822*** (0.06)	0.388*** (0.11)
(f) Temptation \times Bulk (=1)				0.423*** (0.15)
(g) Temptation \times Highest quarter levels only (=1)				0.952*** (0.24)
(h) Temptation \times Bulk \times Highest qrtr. level only (=1)				-0.488 (0.38)
(i) Temptation \times Highest quarter %age only (=1)				-0.482** (0.22)
(j) Temptation \times Bulk \times Highest qrtr. %age only (=1)				0.187 (0.33)
(k) Temptation \times Highest quarter, both (=1)				0.335 (0.28)
(l) Temptation \times Bulk \times Highest quarter, both (=1)				0.401 (0.38)
Observations	11068	11068	11068	11068
R^2	0.43	0.43	0.34	0.35
Household fixed effects	Yes	Yes	Yes	Yes
Item fixed effects	Yes	Yes	No	No
F test p-value: (a)+(b)=0		0.13		0.76
F test p-value: (a)+(c)=0		0.24		0.70
F test p-value: (a)+(d)=0		0.00		0.00
F test p-value: (a)+(f)=0				0.06
F test p-value: (f)+(h)=0				0.85
F test p-value: (f)+(j)=0				0.04
F test p-value: (f)+(l)=0				0.02
F test p-value: (e)+(g)=0				0.00
F test p-value: (e)+(i)=0				0.62
F test p-value: (e)+(k)=0				0.00
F test p-value: (a)+(b)+(f)+(h)=0				0.97
F test p-value: (a)+(c)+(f)+(j)=0				0.04
F test p-value: (a)+(d)+(f)+(l)=0				0.00

Notes: Authors' calculations from SHWALITA data. Standard errors in parentheses; standard errors clustered at household level; *** sig. at 0.01, ** sig. at 0.05, * sig. at 0.1.

retail consumers (Chetty, Looney and Kroft, 2009). Of course, it is possible that inattention combines with other factors to drive households toward small quantity purchasing. We

consider this in Section 6.

5.2.2 Rationing consumption

Could it be that people avoid buying in bulk as a way to limit their consumption? For instance, a sophisticated but present-biased agent would forego bulk purchasing in order to prevent her future self from over-consuming (Laibson, 1997; O’Donoghue and Rabin, 1999). Relatedly, the person or persons who buy food for the household may not trust other household members to control their consumption, and so may limit stocks as a form of rationing. In focus group discussions, we heard variations on this idea, one from a respondent in Bukoba that can be paraphrased as follows: “We know that we need 1 kilogram of maize flour each evening. But if we buy a 50 kilogram bag of maize, we may find that it is gone at the end of one month, because we use too much each day. So it is better to buy smaller amounts.”

If rationing is present, it is most likely to occur on items that are “temptation goods” – goods which, if held in stock, are likely to be over-consumed or consumed too quickly relative to one’s *ex ante* plan. Because temptation goods are culturally specific, we conducted a short survey to rank the study items based on their degree of temptation. We invited 86 Tanzanian field staff members from recent research projects in the country – a panel of experts on household decision-making around economic issues – to rank each of the 21 consumer items in the study on a five-point categorical scale from 1 (not at all tempting) to 5 (tempting for essentially everyone who consumes the item).¹³ Respondents were asked to answer for a typical household in a typical village, not to self-assess their own temptations. The survey was conducted online in June-July 2016. We received 43 responses. We assign each item its average score on the 5-point scale, and then refer to the top third (7 items) as the temptation goods in the study. These are: sugar, rice, cooking oil, soap, tomatoes, cigarettes, and sweet

¹³We excluded the one service in the study – milling of grains – because asking people to characterize milling as tempting or not made no sense during piloting.

bananas.¹⁴

It is worth noting that there is very little correlation between a good having bulk discounts and being classified as tempting. Across the 146 item-district pairs, the correlation between the temptation dummy and the indicator for bulk discounts is -0.07 . The relationship is essentially unchanged if we control for district fixed effects.

We first investigate whether consumption rationing is in fact happening. As a simple test, we ask whether people are prone to make more transactions for temptation goods, consistent with a desire to not accumulate stocks at home. Specifically, we regress K_{hi} on whether an item is a temptation good while controlling for household fixed effects.¹⁵ In column 3 of Table XV, we see that temptation seems to really matter, as households on average make 0.82 more transactions for a temptation good. In other words, households indeed seem to engage in consumption rationing for temptation goods.

We next investigate whether this consumption rationing is driving losses. Specifically, we ask whether there is more consumption rationing for the loss-prone households, which would indicate that the impact of a good being tempting is larger for that group. We also assess whether the relationship between transactions and the existence of bulk discounts (as studied in Section 5.2.1) differs for temptation goods. We are interested in whether the tendency of loss-prone households not to react to bulk discounts could in part be driven by temptation.

Column 4 of Table 15 presents the results. As in column 2, the non-loss-prone households exhibit a strong negative relationship between the number of transactions and the existence of bulk discounts (coefficient (a)), but note that this relationship is weaker for temptation goods (coefficients (a)+(f)). Households with low losses are attending to bulk discounts, but also adjusting their behavior to accommodate temptation.

¹⁴The full ranking and average scores from the temptation survey are shown in online appendix Table A.3.

¹⁵We cannot also control for item fixed effects in these regressions. However, if we exclude the temptation dummy variable and include item fixed effects, coefficient estimates are largely unchanged, indicating that “temptingness” is not conditionally correlated with other unobserved item characteristics that explain variation in K_{hi} . Likewise, the results in columns 1 and 2 of Table XV are qualitatively unchanged if we exclude item fixed effects.

Relative to the other groups, loss-prone households in group (i) (in levels only) demonstrate a strong relationship between temptation and the number of transactions (coefficients (e)+(g)). At the same time, temptation does not alter the relationship between the number of transactions and the existence of bulk discounts for this group. In particular, for both non-temptation goods and temptation goods, this group's purchase patterns are the same whether or not there are bulk discounts (as seen by the fact that coefficients (a)+(b) and coefficients (a)+(b)+(f)+(h) both sum to roughly zero). This pattern of results is consistent with these households being inattentive to bulk discounts while making smaller more frequent purchases for temptation goods. Together, this yields losses on temptation goods.

Loss-prone households in group (ii) (in percentages only) are in fact not influenced by temptation (coefficients (e)+(i)). We know from above that these households are poor, and shop far less than the other groups. While these households are not influenced by bulk discounts for non-tempting goods, for tempting goods they have more transactions when there are bulk discounts (coefficients (a)+(c)+(f)+(j)). That is, they exhibit the surprising pattern that we saw for group (iii) in the previous subsection (in column 2).

Finally, loss-prone households in group (iii) (in both levels and percentages) are influenced mildly by temptation (coefficients (e)+(k)). Moreover, they again exhibit the surprising pattern of having more transactions when there are bulk discounts (coefficients (a)+(d)), and this pattern is even stronger for temptation goods (coefficients (a)+(d)+(f)+(l)).

Overall, these results suggest that consumption rationing might play a role in generating losses. The evidence strongly suggests that consumption rationing is occurring for temptation goods. Moreover, the loss-prone households are inattentive to the financial consequences of this consumption rationing—and for groups (ii) and (iii) bulk discounts are associated with even more transactions for temptation goods. All that said, it is not the case that losses are concentrated on the temptation goods. For the three loss-prone types, losses on the temptation goods are only 38%, 33%, and 41% of total losses. Hence, consumption rationing can be at most part of the story.

5.2.3 Social taxation

For many communities in sub-Saharan Africa, requests by family members and friends for gifts, shared meals, or loans – to which we refer broadly as “social taxes” – are an important part of life (Platteau, 2014). Recent experimental work has shown that participants’ willingness to share windfall gains with others is related in part to the visibility of those gains, suggesting social pressure in favor of redistribution (Goldberg, 2016; Jakiela and Ozier, 2016). Similarly, Baland, Guirkinger and Mali (2011) show in a sample from Cameroon that nearly 1 in 5 people who take out a loan from a savings group do not technically need that loan. These borrowers have adequate cash savings to cover the loan value, but they take out an interest-bearing loan to send a credible signal of poverty to their friends and family, in order to deter requests for handouts. Likewise, De Weerdt, Genicot and Mesnard (2015) show that a transfer recipient’s *perception* of a donor’s wealth affects the value of the transfer between them, conditional on the donor’s actual wealth. This indicates that pressure from transfer recipients is an important factor in determining patterns of social insurance and redistribution.¹⁶

In light of this evidence, a natural hypothesis is that a household might choose not to buy in bulk, despite financial losses, in order to avoid paying social taxes on their store of goods. In support of this mechanism, we present evidence on three points. First, we show that households are indeed subject to social taxes. Second, we show that the social tax rate is higher for those who buy in bulk. Finally, we assess whether households respond to this by reducing their purchase quantities and foregoing bulk discounts. For the final point we rely on anecdotal evidence.

On the first point, the SHWALITA data allow us to directly observe some components of social taxes. In addition to the purchase transactions records, diary keepers also recorded

¹⁶Because the asset in question is a stock of food that can be immediately converted to a consumption good, our measure of social taxes could also be interpreted as a measure of consumption coinsurance (Townsend, 1994; Ligon, 1998; Jalan and Ravallion, 1999; De Weerdt and Dercon, 2006; Kinnan, 2011; Chandrasekhar, Kinnan and Larreguy, 2014). We frame the issue as one of taxes because the behavior of interest is the decision to not buy in bulk, i.e., to forego an investment.

the item description, quantity, unit, and value of any item sold or given away from the household stores. To estimate the household-level social tax rate, we divide the total value of outgoing resources by the total value of incoming resources for the two-week study period. This is not an exact proxy for the household’s social tax rate; the outgoing data include a small number of sales, which may inflate our measure of social tax rates, and measurement error could induce bias in either direction. Yet, to our knowledge this is the first measure of social taxes in the literature that is based on diary records of resource flows into and out of the household.

Table XVI: Social taxation, descriptive statistics

Category	Mean	s.d.
<u>Total value outgoing (TZS)</u>	7982	22567
outgoing: meals and snacks	3028	9526
outgoing: grains	1488	6330
outgoing: pulses and nuts	1414	7449
outgoing: starches	413	2529
outgoing: meat and dairy	260	1406
outgoing: fruits and vegetables	83	769
outgoing: other	1296	10175
<u>Total value incoming (TZS)</u>	72941	101056
incoming: purchases	51184	95335
incoming: own production	15322	24008
incoming: other	6435	11058
Implied social tax rate (%)	12.4	26.1

Notes: Authors’ calculations from SHWALITA data. Estimates based on all relevant activity listed by the 1,499 diary households for all items, not just the 22 items under consideration in the rest of the paper. Figures are the total TZS values of each outgoing or incoming transaction reported in transaction diaries for the categories listed, aggregated to the household-category level by the authors.

Table XVI shows descriptive statistics for the components of incoming and outgoing resources. The uppermost panel shows the value of resources outgoing in the form of sales or gifts, divided into sub-categories. The most important sub-category is “meals and snacks”, which accounts for 40% of outgoing resource flows. Almost all of these are described as a

“full meal” – a guest at the household table (this level of detail not shown). The lower part of the table shows the value of incoming resources. Because we are using these data to characterize households rather than to study purchase behavior, we include all items in the incoming and outgoing measures, not just the 22 items under consideration in the rest of the paper. The mean tax rate is 12.4%, and the median is 2.9%. These descriptive data provide evidence on the first point: social taxes plainly exist and are a common feature of life, although there is significant heterogeneity between households in the social tax rate.

We next ask whether buying in bulk is associated with a higher social tax rate. As a proxy for (not) buying in bulk, we use a household’s total losses L_h , where a smaller L_h is associated with more bulk buying. For robustness, we also estimate models with percentage losses, \tilde{L}_h , as the proxy for not buying in bulk. Table XVII shows the results of regressions of the household social tax rate on L_h or \tilde{L}_h , with and without controls for wealth, demographics, location, and human capital.¹⁷ The coefficients of interest, on L_h and \tilde{L}_h , are stable and highly statistically significant across specifications. Because losses are *decreasing* in bulk purchasing, the negative signs are consistent with the proposed hypothesis: conditional on other characteristics, households that buy in bulk pay higher social tax rates. The magnitudes are economically significant: a one standard deviation increase in L_h (about 935 TZS for the estimation sample) is associated with a tax rate change of -2.06 percentage points, or 17% of the mean. Likewise, a one standard deviation increase in \tilde{L}_h (about 0.066) is associated with a tax rate change of -1.46 percentage points, or 12% of the mean.

Of course, the results in Table XVII merely establish correlation, not that buying in bulk causes households to pay higher social taxes. It is possible that the households that pay higher social taxes for unobserved reasons buy in bulk in order to accommodate the expected requests. However, such reverse causation would contradict the evidence discussed above showing that social tax rates respond to changes in visible resources. Moreover, a second key finding in Table XVII is that social tax rates do not meaningfully co-move

¹⁷Results are similar if we use a Tobit estimator to account for the roughly 35% of households with an estimated zero tax rate. See online appendix Table A.4.

Table XVII: Household-level social tax rate regressed on losses and household characteristics

Dependent variable: household-level social tax rate				
	(1)	(2)	(3)	(4)
L_h	-0.0021*** (0.0007)	-0.0022*** (0.0008)		
\tilde{L}_h			-20.6654** (9.0284)	-22.0560** (9.3974)
Wealth index quartile 2 (=1)		1.7832 (2.1914)		1.3782 (2.1465)
Wealth index quartile 3 (=1)		2.4235 (2.0651)		1.8586 (2.0382)
Wealth index quartile 4 (=1)		2.0468 (2.9586)		0.7097 (2.8818)
Household size		0.0079 (0.2715)		-0.1438 (0.2720)
Age of head (years)		-0.0118 (0.0506)		-0.0099 (0.0504)
Head is female (=1)		1.4034 (1.7317)		1.3932 (1.7450)
Distance to center of community (km)		-0.2208 (0.7283)		-0.1607 (0.7189)
Observations	1472	1465	1472	1465
R^2	0.09	0.09	0.09	0.09

Notes: Authors' calculations from SHWALITA data. Standard errors in parentheses; standard errors clustered at village level; *** sig. at 0.01, ** sig. at 0.05, * sig. at 0.1. All regressions include district fixed effects and controls for questionnaire module.

with important observable household characteristics, such as wealth, size, and location. We find this surprising. Yet, this result lends support to the idea that social tax rates are not exclusively determined by time invariant household characteristics, but respond to observable resources such as the household stocks on hand.

Finally, we assess whether fear of paying social taxes might be the reason for incurring losses by not buying in bulk. On this point we rely on anecdotal evidence from our qualitative work. The case was stated most eloquently by someone we interviewed in the Kagera region: “If I buy 5 kilograms of sugar, everyone will take their tea at my house.” We heard variations on this idea from other people. Buying small quantities – or, at least, not buying very large quantities – can be a useful way to deter requests from one’s social network. Even if goods can be stored privately so that explicit redistributive pressure is not a problem, people may hold internalized norms in favor of sharing that which is available, which prevents them from

stocking too much of an item. The end result is that households incur losses on some items for which they would otherwise pay a lower unit price.

This analysis raises the question of whether small quantity purchasing to avoid social taxation is financially inefficient. We could attempt to answer this question by calculating the counterfactual tax rate that households would pay if they purchased in bulk, and refining our measure of L_h accordingly. We opted not to do this, because construction of such estimates requires a parametric model relating social taxes to purchase quantities and other household characteristics. If social tax rates are determined in part by unobserved household characteristics, as they likely are, then such a model would be of little use for the out-of-sample predictions required by the exercise. Given this inherent limitation, we prefer to simply emphasize that foregoing bulk purchasing may be financially efficient for some households, because it saves them from paying higher social taxes. We discuss this more broadly in Section 6.

5.2.4 Coordination costs within the household

Finally, we examine the possibility that purchase of financially inefficient small quantities could be driven by the challenge of coordinating purchases between household members. We consider two types of coordination problems: the first related to the division of shopping responsibilities between household members, and the second to the relationship between spouses.

To implement this analysis we need a measure of shopping coordination (or lack thereof). As a proxy for uncoordinated shopping, we use the number of days on which two or more household members purchase the same item. This measures the key behavior of interest, because households that do not coordinate the shopping of their members incur losses when two or more people purchase quantities of an item that would have cost less, total, if purchased all at once. For these tests we use data from the 500 households that were randomly assigned to the “personal diary” treatment. Each adult in these households

completed his or her own transaction diary, allowing us to observe which individual was responsible for each purchase. The modal household has 2 personal diary keepers (54% of observations); the mean number of diary keepers is 2.1, and the maximum is 7. For almost 40% of personal diary households, and for half of those with more than one diary keeper, there is at least one instance of two or more members buying the same good on the same day. Across all households, the mean number of days with multiple purchasers of the same item is 1.06.

Table XVIII: Loss regressed on number of days with multiple purchasers of same item

Households:	All	Multi-diary	All	Multi-diary
Dependent variable:	Loss	Loss	% Loss	% Loss
	(1)	(2)	(3)	(4)
Number of days 2+ people buy same item	67.505** (28.26)	63.789** (29.48)	-0.001 (0.00)	-0.001 (0.00)
Household size	42.287* (21.87)	28.907 (26.74)	-0.001 (0.00)	-0.002 (0.00)
Wealth index quartile 2 (=1)	90.348 (114.58)	87.828 (135.05)	-0.005 (0.01)	-0.003 (0.01)
Wealth index quartile 3 (=1)	74.841 (133.06)	70.731 (152.44)	-0.007 (0.01)	-0.007 (0.01)
Wealth index quartile 4 (=1)	183.368 (177.12)	73.947 (211.43)	-0.019 (0.01)	-0.014 (0.01)
Observations	495	381	495	381
R^2	0.18	0.19	0.14	0.16

Notes: Authors' calculations from SHWALITA data. Standard errors in parentheses; standard errors clustered at village level; *** sig. at 0.01, ** sig. at 0.05, * sig. at 0.1. All regressions include district fixed effects. Personal diaries only.

Table XVIII shows results of regressing losses at the household level on the multiple-purchaser variable and key control variables. Results are shown for all households in the personal diary treatment as well as only for those households with multiple diary-keepers, and for both level and percentage losses. In column 1 we see that the level value of losses is increasing in the degree of un-coordinated shopping. Each additional day on which more than one person buys the same item is associated with 68 TZS higher losses. This result is essentially unchanged if we restrict attention to households with multiple diary-keepers (column 2). Importantly, these results hold while conditioning on household size. We expect

more frequent multiple-purchasing in larger households, because the costs of coordination are increasing in the number of people to coordinate. We have already seen that losses are conditionally increasing in household size (Table VIII).

Columns 3 and 4 report estimates of similar regressions with \tilde{L}_h as the dependent variable. Here we see that there is no link between lack of coordination and percentage losses. This is not surprising, as we already know that households with high percentage losses tend to be smaller and to make many fewer purchases than average (Table IX).

A second type of coordination problem was mentioned by numerous respondents during our qualitative work. The suggestion was that husbands in Tanzania ration the spending of their wives by giving them daily allowances to purchase necessities, such as the components of the family meal. The husband's goals, according to interviewees, are to carefully guard the household finances and to prevent the wife from spending money on items for herself. This mechanism alone would not be sufficient to generate losses, because the wife could potentially save some cash and delay purchasing certain items in order to buy goods in bulk. However, if such behavior would be perceived as a violation of the social contract between spouses, the personal cost to the woman could be too high to justify saving up in order to bulk purchase.

This model of spousal control is difficult to test, because there are many unobserved forces that could lead to gender differences in purchasing behavior. We can, however, use the data from personal diaries to examine whether losses are overwhelmingly attributable to women rather than men. For many items women do buy slightly smaller quantities than men, on average. Yet, if we treat each individual as an independent unit and calculate the self-financed purchasing delay for each item that appears in each personal diary (as in Table XIII), the median days to efficient purchasing are 7.9 for women and 9.3 for men. This suggests that financial losses are not attributable to small quantity purchasing by women alone. Thus, while we cannot rule out that husbands impose short-term spending constraints on their wives, we do not have evidence that this leads to the accrual of losses.

6 Discussion

The analysis in this paper has demonstrated the existence of significant bulk discounts in Tanzanian markets, even at relatively low quantities. We have shown that some households suffer substantial financial losses by not taking advantage of those bulk discounts. We have further investigated a number of possible explanations for this behavior, rejecting some mechanisms and presenting evidence in favor of others. In this section, we discuss some broader implications of our findings, and conclude.

A clear implication of our analysis is that consumer behavior seems driven by multiple mechanisms. Indeed, we would argue that our approach of allowing for and testing a range of possible mechanisms yields a more complete picture of household behavior than if we had focused on identifying one particular mechanism. Moreover, there are likely to be some interactions between these mechanisms. Most notably, while inattention by itself can lead to a failure to take advantage of bulk discounts, it is perhaps more important as an enabling mechanism that “multiplies” the impact of other mechanisms. Self-rationing, avoidance of social taxation, and coordination problems all create a clear reason to buy more frequently and in smaller quantities. If a household attends to the impact of bulk discounts, it will weigh the benefits of purchasing in smaller quantities against the financial cost of forgoing bulk discounts. But inattention can render a household particularly prone to losses.

A related point is that there is no reason to believe that a single mechanism or constellation of mechanisms should serve to explain the behavior of all households. Our analysis clearly indicates the existence of heterogeneity in households. At one level, we see significant heterogeneity across households in terms of losses incurred. At a second level, even among the loss-prone households, we see different types—e.g., the households who are in the top quartile of level losses but not in the top quartile of percentage losses appear to be very wealthy households, whereas those who are in the top quartile of percentage losses but not level losses appear to be very poor households. Finally, we see evidence that households are impacted differentially by the various mechanisms—e.g., in Table 15 we see

that the wealthy loss-prone households seem especially influenced by temptation, while the poor loss-prone households seem not to be influenced by temptation.

Each of the four mechanisms implicated by our analysis represents a natural response to some constraint faced by the shopper. Inattention to bulk discounts is a natural response to the cognitive cost of figuring out the financial implications. Limiting stocks is a natural response to avoid inefficient over-consumption by one's future self or by other household members, and to the social taxes that emerge from longstanding traditions of meal-sharing and coinsurance in village communities. And limiting purchases to avoid excessive total household purchases is a natural response to coordination issues. Of course, the big question is whether households are responding optimally to these other constraints. Our data do not permit us to explicitly address this question. Yet, given the magnitude of the financial costs, we find it hard to believe that this behavior could be fully optimal, especially for inattention.

Our finding that households in Tanzania appear to forego potential gains in consumption by not buying in bulk should clearly be of interest to policymakers and consumer advocates. Given our analysis, are there policy interventions that might make consumers better off by helping them to take more advantage of bulk discounts? The answer depends on the mechanism. Perhaps most promising would be interventions designed to make bulk discounts more salient, or otherwise help households to better understand the full financial implications of not buying in bulk. But one could also imagine interventions designed to help people better manage their consumption of stocks, both within the household and with neighbors, or interventions designed to better manage the coordination of household purchases. Indeed, on the last point, we wonder whether the dramatic increase in cell phone use in Tanzania since the study period may have already had a significant impact by lowering the costs of coordination. For all such interventions, however, one must not ignore that reductions in small-quantity purchasing could lead to negative supply-side responses and resulting general equilibrium effects (which we cannot investigate with these data).

Finally, we conclude by reiterating our general contribution of describing and quan-

tifying from a new angle an important feature of life in a low-income country. Spending time in villages in Tanzania or similar countries, one cannot help but notice the frequency of small quantity transactions at kiosks, shops, markets, and roadside stands. We examine the financial implications of this way of making purchases, and provide evidence on numerous hypotheses for why people might arrange their spending this way. We find strong evidence against the "poor pay more" hypothesis, calling into question some current wisdom about how inequities may persist or be exacerbated by differential access to bulk discounts. We also find a combination of external and internal forces that rationalize the observed purchasing patterns. In this regard, we believe that our paper makes a novel contribution toward understanding the micro-foundations of consumer choice in developing economies.

A Appendix — for online publication only

Table A.1: Regressions of expenditure on quantity and its square

Item	Coefficient on q^2	
	point estimate	standard error
Rice	-1.33	(4.47)
Maize	-1.22	(0.73)
Flour	-22.14**	(6.11)
Milling	-0.06	(0.23)
Cassava	0.31	(2.98)
Cooking Bananas	-2.83	(1.81)
Sugar	-16.61	(8.87)
Beans	-0.10	(7.51)
Coconut	-14.58	(18.14)
Tomatoes	-69.10	(42.87)
Onions	-60.72	(51.09)
Sweet Bananas	-78.50**	(26.41)
Sardines	-237.81***	(60.52)
Cooking Oil	-407.26***	(109.55)
Salt	-17.32	(33.65)
Tea	-13.19**	(4.53)
Tea Leaves	-6806.18***	(803.50)
Charcoal	-9.26**	(2.68)
Kerosene	-240.55***	(47.57)
Matches	-1.03***	(0.18)
Soap	-4.92***	(0.52)
Cigarettes	-0.06	(0.03)

Notes: Authors' calculations from SHWALITA data. Standard errors in parentheses; standard errors clustered at district level; *** sig. at 0.01, ** sig. at 0.05, * sig. at 0.1. Each coefficient is from a separate regression of transaction-level expenditure on quantity and quantity squared, for only the item indicated. We report the coefficient on quantity squared.

Table A.2: Counterfactual change in number of transactions, if purchasing at lowest unit price

Subgroup:	25% highest loss, level only			25% highest loss, %age only			25% highest loss, both		
	Mean	Mean	Mean	Mean	Mean	Mean	Mean	Mean	Mean
Item	K_{hi}	K_{hi}^*	diff.	K_{hi}	K_{hi}^*	diff.	K_{hi}	K_{hi}^*	diff.
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Kerosene	5.5	1.8	-3.7	3.0	0.9	-2.1	4.7	1.3	-3.4
Onions	5.6	4.8	-0.8	2.3	2.3	0.0	4.5	4.1	-0.4
Cooking Bananas	3.0	2.2	-0.8	1.9	1.0	-0.9	3.7	2.2	-1.5
Cooking Oil	7.1	4.6	-2.6	3.2	2.0	-1.3	6.6	3.0	-3.6
Tea Leaves	5.1	4.9	-0.2	1.5	0.7	-0.8	3.5	1.5	-2.0
Sardines	3.8	5.6	1.8	3.3	3.8	0.5	4.2	5.1	0.9
Tomatoes	7.1	8.0	0.9	3.2	3.5	0.4	6.2	6.3	0.1
Salt	2.4	2.2	-0.1	1.8	1.5	-0.3	2.3	1.9	-0.4
Coconut	6.1	4.7	-1.4	2.0	1.5	-0.5	3.8	2.5	-1.3
Maize	3.2	2.7	-0.6	1.3	1.6	0.3	3.5	2.0	-1.5
Sweet Bananas	3.0	5.0	2.0	1.4	2.0	0.6	2.4	4.2	1.8
Cassava	2.3	4.4	2.1	2.0	4.5	2.6	1.9	5.1	3.2
Soap	3.9	5.9	2.0	2.3	2.8	0.5	3.3	4.9	1.6
Charcoal	7.5	9.9	2.4	3.0	1.9	-1.1	5.9	5.9	0.1
Milling	3.1	4.7	1.6	1.9	2.6	0.6	3.4	3.8	0.4
Matches	2.4	3.5	1.1	1.8	1.8	-0.0	2.3	3.3	1.0
Cigarettes	7.1	29.4	22.3	3.5	6.1	2.6	5.9	15.1	9.2
Sugar	5.6	9.9	4.3	2.0	3.6	1.6	3.7	6.0	2.3
Flour	7.4	17.3	9.9	2.5	6.4	3.9	4.1	10.7	6.6
Rice	4.8	12.0	7.2	1.4	2.7	1.4	2.9	5.2	2.3
Beans	3.2	7.7	4.5	1.4	3.5	2.1	2.5	5.0	2.5
Tea	4.9	7.3	2.4	1.8	3.1	1.3	3.8	6.4	2.6
AVERAGE	5.4	8.0	2.6	2.4	2.6	0.2	4.0	4.3	0.4

Notes: Authors' calculations from SHWALITA data. Items ordered by decreasing values of mean \tilde{L}_h , from column 6 in Table VII.

Table A.3: Temptation survey results (1 = Not tempting; 5 = Highly tempting)

Item	Mean score
Tea leaves	2.31
Maize	2.33
Cassava	2.43
Kerosene	2.54
Sardines	2.74
Salt	2.79
Matches	2.81
Coconut	2.88
Cooking bananas	3.00
Onions	3.00
Flour	3.02
Beans	3.02
Prepared tea	3.15
Charcoal	3.19
Sweet bananas	3.38
Cigarettes	3.81
Soap	3.83
Tomato	3.83
Cooking oil	3.85
Rice	4.29
Sugar	4.31

Notes: Authors' calculations from survey conducted with 43 Tanzanians.

Table A.4: Household-level social tax rate regressed on losses, Tobit

Dependent variable: Estimated household-level social tax rate				
	(1)	(2)	(3)	(4)
L_h	-0.0022** (0.0011)	-0.0024** (0.0012)		
\tilde{L}_h			-26.7382* (14.1943)	-25.8822* (14.5389)
Wealth index quartile 2 (=1)		1.1341 (2.6910)		0.6754 (2.6420)
Wealth index quartile 3 (=1)		2.5567 (2.5476)		1.9555 (2.5404)
Wealth index quartile 4 (=1)		1.4114 (3.8607)		-0.0130 (3.7960)
Household size		0.4074 (0.3631)		0.2495 (0.3622)
Age of head (years)		-0.0464 (0.0725)		-0.0440 (0.0721)
Head is female (=1)		0.5871 (2.5705)		0.5962 (2.5959)
Distance to center of community (km)		0.4028 (1.0403)		0.4775 (1.0360)
Observations	1472	1465	1472	1465
R^2	0.03	0.03	0.03	0.03

Notes: Authors' calculations from SHWALITA data. Standard errors in parentheses; standard errors clustered at village level; *** sig. at 0.01, ** sig. at 0.05, * sig. at 0.1. All regressions include district fixed effects and controls for questionnaire module.

References

- Alby, Philippe, Emmanuelle Auriol, and Pierre Nguimkeu.** 2013. “Social barriers to entrepreneurship in Africa: The forced mutual help hypothesis.” Working paper.
- Anderson, Siwan, and Jean-Marie Baland.** 2002. “The Economics of Roscas and Intra-household Resource Allocation.” *The Quarterly Journal of Economics*, 117(3): 963–995.
- Ashraf, Nava, Dean Karlan, and Wesley Yin.** 2006. “Tying Odysseus to the mast: Evidence from a commitment savings product in the Philippines.” *The Quarterly Journal of Economics*, 121(2): 635–672.
- Attanasio, Orazio, and Christine Frayne.** 2006. “Do the poor pay more?” Working paper.
- Attanasio, Orazio, and Elena Pastorino.** 2015. “Nonlinear Pricing in Village Economies.” NBER working paper 21718.
- Baland, Jean-Marie, Catherine Guirkinger, and Charlotte Mali.** 2011. “Pretending to be poor: Borrowing to escape forced solidarity in Cameroon.” *Economic Development and Cultural Change*, 60(1): 1–16.
- Beegle, Kathleen, Joachim De Weerd, Jed Friedman, and John Gibson.** 2012. “Methods of household consumption measurement through surveys: Experimental results from Tanzania.” *Journal of Development Economics*, 98(1): 3–18.
- Burke, Marshall.** 2014. “Selling low and buying high: An arbitrage puzzle in Kenyan villages.” Working Paper.
- Chandrasekhar, Arun G, Cynthia Kinnan, and Horacio Larreguy.** 2014. “Social networks as contract enforcement: Evidence from a lab experiment in the field.” NBER working paper 20259.

- Chetty, Raj, Adam Looney, and Kory Kroft.** 2009. “Salience and Taxation: Theory and Evidence.” *The American Economic Review*, 99(4): 1145.
- Deaton, Angus.** 1988. “Quality, quantity, and spatial variation of price.” *The American Economic Review*, 78(3): 418–430.
- Deaton, Angus, Jed Friedman, and Vivi Alatas.** 2004. “Purchasing power parity exchange rates from household survey data: India and Indonesia.” Princeton Research Program in Development Studies Working Paper.
- De Weerdt, Joachim, and Stefan Dercon.** 2006. “Risk-sharing networks and insurance against illness.” *Journal of Development Economics*, 81(2): 337–356.
- De Weerdt, Joachim, Garance Genicot, and Alice Mesnard.** 2015. “Asymmetry of information within family networks.” National Bureau of Economic Research working paper 21685.
- Duflo, Esther, Michael Kremer, and Jonathan Robinson.** 2011. “Nudging farmers to use fertilizer: Theory and experimental evidence from Kenya.” *The American Economic Review*, 101(6): 2350–2390.
- Dupas, Pascaline, and Jonathan Robinson.** 2013. “Savings constraints and microenterprise development: Evidence from a field experiment in Kenya.” *American Economic Journal: Applied Economics*, 5(1): 163–192.
- Filmer, Deon, and Lant H Pritchett.** 2001. “Estimating wealth effects without expenditure Data—Or tears: An application to educational enrollments in states of India.” *Demography*, 38(1): 115–132.
- Gibson, John, and Bonggeun Kim.** 2011. “A Simple Test of Liquidity Constraints and Whether the Poor Pay More.” Working paper.

- Goldberg, Jessica.** 2016. “The effect of social pressure on expenditures in Malawi.” Working paper.
- Griffith, Rachel, Ephraim Leibtag, Andrew Leicester, and Aviv Nevo.** 2009. “Consumer shopping behavior: How much do consumers save?” *The Journal of Economic Perspectives*, 23(2): 99–120.
- Jack, B Kelsey, and Grant Smith.** 2015. “Pay as You Go: Prepaid Metering and Electricity Expenditures in South Africa.” *The American Economic Review: Papers and Proceedings*, 105(5): 237–241.
- Jakiela, Pamela, and Owen Ozier.** 2016. “Does Africa need a rotten kin theorem? Experimental evidence from village economies.” *Review of Economic Studies*, 83(1): 231–268.
- Jalan, Jyotsna, and Martin Ravallion.** 1999. “Are the poor less well insured? Evidence on vulnerability to income risk in rural China.” *Journal of Development Economics*, 58(1): 61–81.
- Kinnan, Cynthia.** 2011. “Distinguishing barriers to insurance in Thai villages.” Working paper.
- Kremer, Michael, Jean Lee, Jonathan Robinson, and Olga Rostapshova.** 2013. “Behavioral biases and firm behavior: Evidence from Kenyan retail shops.” *The American Economic Review: Papers and Proceedings*, 103(3): 362–368.
- Kremer, Michael, Jean Lee, Jonathan Robinson, and Olga Rostapshova.** 2015. “Rates of Return, Optimization Failures, and Behavioral Biases: Evidence from Kenyan Retail Shops.” Working paper.
- Laibson, David.** 1997. “Golden eggs and hyperbolic discounting.” *The Quarterly Journal of Economics*, 112(2): 443–477.

- Ligon, Ethan.** 1998. "Risk sharing and information in village economies." *The Review of Economic Studies*, 65(4): 847–864.
- Mani, Anandi, Sendhil Mullainathan, Eldar Shafir, and Jiaying Zhao.** 2013. "Poverty impedes cognitive function." *Science*, 341(6149): 976–980.
- McKelvey, Christopher.** 2011. "Price, unit value, and quality demanded." *Journal of Development Economics*, 95(2): 157–169.
- Mullainathan, Sendhil, and Eldar Shafir.** 2013. *Scarcity: Why Having Too Little Means so Much*. Macmillan.
- Mussa, Richard.** 2015. "Do the Poor Pay More for Maize in Malawi?" *Journal of International Development*, 27(4): 546–563.
- O'Donoghue, Ted, and Matthew Rabin.** 1999. "Doing it now or later." *American Economic Review*, 89(1): 103–124.
- Platteau, Jean-Philippe.** 2006. "Solidarity norms and institutions in village societies: Static and dynamic considerations." *Handbook of the Economics of Giving, Altruism and Reciprocity*, 1: 819–886.
- Platteau, Jean-Philippe.** 2014. "Redistributive pressures in Sub-Saharan Africa: Causes, consequences, and coping strategies." *Africa's Development in Historical Perspective*, 153.
- Rao, Vijayendra.** 2000. "Price heterogeneity and "Real" inequality: A case study of prices and poverty in rural south India." *Review of Income and Wealth*, 46(2): 201–211.
- Sahn, David E., and David Stifel.** 2003. "Exploring Alternative Measures of Welfare in the Absence of Expenditure Data." *Review of Income and Wealth*, 49(4): 463–489.
- Stephens, Emma C, and Christopher B Barrett.** 2011. "Incomplete credit markets and commodity marketing behaviour." *Journal of Agricultural Economics*, 62(1): 1–24.

Subramanian, Shankar, and Angus Deaton. 1996. “The demand for food and calories.”
Journal of Political Economy, 104(1): 133–162.

Townsend, Robert M. 1994. “Risk and insurance in village India.” *Econometrica*,
62(3): 539–591.

Zeitlin, Andrew. 2011. “Understanding heterogeneity: Risk and learning in the adoption
of agricultural technologies.” Georgetown University working paper.