“Sell Low, Buy Low?'' - A New Explanation for a Persistent Puzzle

Lila Cardell, University of Illinois at Urbana-Champaign, Lilac2@illinois.edu
Hope Michelson, University of Illinois at Urbana-Champaign, hopecm@illinois.edu

Selected Paper prepared for presentation at the 2020 Agricultural & Applied Economics Association Annual Meeting, Kansas City, MO
July 26-28, 2020

"Sell Low, Buy Low?" - A New Explanation for a Persistent Puzzle

Lila Cardell and Hope Michelson*

June 28, 2020

Abstract

We propose a new explanation for the commonly-observed "sell low, buy high" behavior among small-scale staple grain farmers in developing countries: risk aversion, combined with fact-based expectations of negative returns from delayed sales, incentivizes farmers to opt out of storage. Using 20 years of data from 751 markets in 26 countries in Sub-Saharan Africa, we demonstrate that the market prices in lean seasons (the assumed "higher price" seasons) often fail to rise above prevailing prices in the harvest seasons (the "low price" seasons), and that the probability of *negative* returns to storage across seasons is 28.6%. Our results indicate that storing does not stochastically dominate immediate post-harvest sales in any country, given the substantial probability of negative relative returns. We use a two-period model to demonstrate that aversion to negative returns can induce households that both produce and consume staple cereals to select out of staple grain storage, even when credit and storage options are available.

*Cardell: University of Illinios at Urbana-Champaign, lilac2@illinois.edu. Michelson: University of Illinios at Urbana-Champaign, hopecm@illinois.edu

Staple cereal grain prices exhibit recurring patterns of seasonal price fluctuations in rural markets in developing countries, with low prices at harvest, followed by steady rises to an annual high shortly before lean-season planting (Kaminski et al., 2016, 2014). This pattern is a documented contributor to seasonal hunger, malnutrition, and food insecurity among small farmer households (Sahn, 1989; Christian and Dillon, 2018).

A contributing factor and a persistent puzzle: small farm households have often proved unable or unwilling to exploit attractive inter-temporal arbitrage opportunities for storable commodities (maize, rice). Instead, they are observed to sell their output when prices are low, often repurchasing the same commodity for family sustenance later in the year when prices tend to be substantially higher (Barrett, 2007). As Stephens and Barrett (2011) observe, farmers engaging in this pattern of selling low and buying high are effectively using staple grain markets as a lender of last resort. Explanations for this pattern of behavior have centered on liquidity constraints and transaction costs (Stephens and Barrett, 2011; Burke et al., 2019; Aggarwal et al., 2018) as well as inadequate storage technologies (Walker et al., 2018; Channa et al., 2019).

We provide a new and better explanation for this documented reluctance among small farmers to store across seasons: a perceived risk, founded in prior experience, of realizing negative returns on stored cereals when lean season prices fail to exceed the harvest price. In this scenario, credit availability and storage access do not affect storage uptake; price risk would be reduced only through forward or future contracts that lock in a price to prevent negative returns.

We use data from 751 markets in 26 countries to establish the distribution of seasonal price increases and analyze the risk-management implications for small farmers. We demonstrate that market prices in the lean season (the "high price" season) do not always exceed the prevailing price during the harvest season; moreover, we find that the probability of *negative* returns to storage is 28.6% across markets and years, ranging from a low of 10.5% in Mozambique to a high of 55.6% in Nigeria. In fact, a lean season price can not only fail to rise relative to the harvest season price; it can also prove considerably lower. For example, when the lean season price fails to exceed the harvest price range, the average difference has recently ranged between 4.5% lower in Burkina Faso and 34.3% lower in South Sudan. Our comparisons are conservative, abstracting away from the costs of storage equipment and space, post-harvest losses in storage, foregone interest on sales revenue, and any added costs farmers incur associated with selling in the lean season.

For any market series where there is a substantial probability of prices falling post-harvest, storing to sell later may not stochastically dominate for households with utility that is monotonically increasing and concave in income. Households that both produce and consume maize face two sources of risk from stochastic prices: income risk and consumption price risk. For net seller households the former dominates; and uncertainty causes risk-averse households to under-invest in profit-generating opportunities, e.g. low storage uptake. For net buyer households, the latter dominates and uncertainty causes risk-averse households to over-invest in profit-generating opportunities, e.g. excess storage. (Finkelshtain and Chalfant, 1991; Barrett, 1996; Saha and Stroud, 1994) Under the assumption of positive returns, both net seller and net buyer households should store, and depending on the available terms, take credit to do so. However, the possibility of negative returns would discourage all net sellers from storing, and potentially net buyer households as well. Credit access would not affect storage decisions for net sellers facing possible negative returns. We focus on net seller households, who have a clear behavioral choice in responding to price risk.

Based on insights from analysis of the price data, we use a simple two-period model to assess the welfare implications of storing. Our results demonstrate that in most countries under review, storing does not always stochastically dominate immediate post-harvest sales in any order of dominance because of a substantial probability of negative relative returns, even when those returns are expected to be much higher. This constitutes a new insight in the literature, helping to explain why farmers opt to sell immediately post-harvest if they have no hedging options, and also why small-scale traders (who lack the capital to engage in spatial arbitrage opportunities) remain unwilling to invest.

Existing studies documenting price fluctuations and proposing interventions related to storage and credit have two important limitations: first, these analyses are commonly based on only one or two years of data; second, they normally analyze the differences between mean harvest season prices and lean season prices by averaging either across years or across markets or both. [1] By using more data over a longer time span, we show patterns that would otherwise be missed, and demonstrate how the variations in price can affect storage choice. A stochastic dominance approach allows us to avoid placing restrictions on consumer preferences, and to focus on the relevance of policies related to price stabilization, namely futures contracts. We use hypothesis tests developed by Abadie (2002); Barrett and Donald (2003); Linton et al. (2005) and others to evaluate the likelihood of stochastic dominance between two distributions.

Economists have long focused on the effects of price volatility and stabilization on consumers and producers (Waugh, 1944; Oi, 1961; Stiglitz, 1969). While financial derivatives such as futures markets offer an efficient approach to smoothing prices across seasons, they are not generally applicable to small-scale farming in low income countries. Numerous interventions by Non-Governmental Organizations and researchers have been designed to provide credit and storage options to farmers, and a wealth of Randomized Control Trials (RCTs) have been implemented in recent years to evaluate such efforts: Burke et al. (2019) provide credit to farmers in Kenya; Basu and Wong (2015) distribute storage equipment to farmers in Indonesia; Aggarwal et al. (2018) encourage communal

[1] Indeed, by pulling individual years of data or by averaging across years and markets, we can replicate the results and graphs suggesting the presence of inter-temporal arbitrage opportunities that drive the existing literature and associated interventions to promote farmer storage across seasons.

maize storage in Kenya; and Channa et al. (2018) combine storage and credit in Tanzania. Our results indicate that the probability of negative returns is an important deterrent to storage by small farmers and small traders, and suggest new directions for research and policy in this critical area.

1 Empirical Analysis

1.1 Data

The World Food Programme (WFP) food price monitoring system reports monthly food prices using data collected by WFP and national agricultural ministries. (Oscar Maria Caccavale and Flämig, 2017) Data is available at a sub-national level for food staples, fruits, vegetables, and animal products. We select all countries in Sub-Saharan Africa (SSA) with monthly retail prices available for maize. If more than one variety of maize was available for a given market, e.g. yellow and white maize, we chose the country's more predominant variety. We convert prices to USD using exchange rates from the International Monetary Fund and adjust for inflation, using January 2010 as a base for all countries and markets.

The analysis requires that we identify the harvest and lean season for each market. Agricultural season data is collected by the Food and Agricultural Organization (FAO) for the Global Information and Early Warning System (GIEWS.) GIEWS reports national and sometimes sub-national harvest and planting season dates for various crops, with data provided by national ministries.

We merge the GIEWS agricultural season designations with the WFP price data to identify the prices for the months of the year that GIEWS reported as "harvest" or "planting". If GIEWS reports multiple agricultural regions within a country, we use the maize season data located closest to the market coordinates, within the same country.

We calculate inter-seasonal price differences (and associated returns) by replicating the "seasons" from the perspective of a farmer considering grain storage at the end of harvest: we create a "harvest season price" as the last price of the months designated by GIEWS as harvest months for a given market. The "lean season price" is the maximum price of the months designated as planting. This combination of prices has the advantage of being quite conservative, and gives us a lower bound on the probability of negative returns. We calculate returns for each season as the percent change from the "harvest season price" to the following "lean season price".

We retain in the data all years for which we have price data for both the harvest season and the following lean season: for example, if the harvest occurs in September-October of 2018 and planting in January-February 2019, the return for the 2018 season would be the difference from the October 2018 "harvest" price to either the January or February 2019 "lean" price, depending on whichever month had the higher price. If either the harvest or planting season price was unavailable for that market-year, we do not include the observation in our data.

GIEWS reports multiple maize seasons for a subset of countries with two growing seasons per year, which we account for by calculating the average return across the two seasons. We remove 69 markets with only one year of data. Our final data include 5922 market-year observations across 751 markets in 26 countries in SSA between the years 2000-2020.

The GIEWS data provides harvest and planting dates for maize crops at a national and occasionally sub-national level, possibly masking relevant regional variation in the timing of harvest and planting. The harvest and planting seasons do not always coincide with the months of the minimum and maximum retail prices observed in a given year. We select prices to represent the farmers' decision-making problem – whether to store or sell at the end of harvest and the return they would have received in the lean season if they chose to store maize in a given year.

1.2 Results: Price differentials across seasons

Using monthly retail price data collected over the last 20 years from 751 maize markets across 26 countries, we find evidence of both positive and negative price differentials between seasons. The phenomenon of negative price differentials across harvest and lean season is widespread, not confined to any country or set of years. This finding is contrary to prior research that has assumed that higher lean season prices ensure positive returns to storing grain at harvest.

In Table 1, we present a summary of the data and findings. For each country, we list the years for which data was available, the number of markets in Column (1) and the total number of market-years in Column (2) (ie: for Benin, there are 64 total market-year observations across 19 markets and nine years.) Column (3) shows the probability of negative returns: the proportion of market-years in each country in which the price decreased from harvest to the following lean season. Column (4) presents the average returns by country across all market-years. Column (5) presents the average returns in each country for market-years where the price increased from harvest to the subsequent lean season (the "sell low, buy high" phenomenon.) Column (6) presents the average negative returns for market-years where the price decreased, i.e. the alternative and less-discussed case: years in which the harvest season price exceeded the lean season price.

The results presented in Table 1 indicate that farmers across Africa face both positive and negative inter-seasonal price spreads (Table 1); years in which the price rises significantly after harvest as well as years in which the price stays flat or even declines in the lean season relative to its level at harvest. Table 1 shows that this phenomenon is not restricted to particular countries.

We discuss a few examples to clarify the findings presented in Table 1. The Malawi data include 762 market-year observations across 71 markets. In 19.7% of the observed market-years, the lean season price is below the harvest season price by an average of 11.6%. For the remaining 80.3% of market-years, the lean season price does exceed the harvest season price, and it exceeds it by 41.8% on average. Over all market-years, the returns were positive: 31.3%. Another example: in Senegal, the lean season price is lower than the harvest season price for 33.1% of the market-years, by an average of 9.5%.

Figure 1 presents three depictions of observed trends. In Figure (1a), we present the distribution of returns for each market-year across years between 2000 and 2020. Figure (1a) demonstrates that the phenomenon is not restricted to particular years, nor is it attenuating in time.

Figure (1b) shows the frequency and intensity of the negative returns phenomenon. The figure demonstrates that even in markets where expected returns are high, the risk of loss is nontrivial. Each dot represents one of the 751 markets, with the percent of seasons when the harvest season price exceeded the lean season on the x-axis and the average returns for that market on the y-axis. Consistent with Table 1 and Figure (1a), returns are generally positive on average for a given market across years. For 125 markets, returns are always positive, and for nine markets, returns are always negative. Of those nine markets, the market-year observations cover eight different countries, and nine different years.

Figure (1c) demonstrates seasonal price trends in three dimensions. Each dot represents a market-year, with the color of the dot representing the seasonal return: green is negative returns and positive returns increase in color saturation with the most purple the most positive. The x and y axis are the z-scores of harvest and lean season prices respectively, with the z-score calculated as the price related to the mean harvest or lean season price for each market.

- Quadrant I shows the "sell low, buy high" phenomenon, with harvest prices lower than usual (possibly due to a good harvest, or higher sales than usual) and lean season prices higher than usual (possibly due to many households buying back the same grain they sold in the harvest season). For farmers who stored at harvest, returns are high in these market-years.

- Quadrant II contains market-years with relatively high harvest and high lean season prices, and varying returns. These could be market-years with bad harvests and varying levels of government intervention to alleviate maize shortages in the lean season.

- Quadrant III shows market-years with higher than average harvest prices and lower than average lean season prices. Points in this quadrant are associated with zero or negative returns. These could be bad harvest years, with government intervention, or lower sales than usual, in which case farmers who store grain experience negative returns in the lean season when prices fall.

- Quadrant IV includes market-years with relatively low harvest and lean season prices, with variable returns as well, potentially representing market-years with good harvests that carry over into lean season abundance.

An additional insight from Figure (1c) is important: the possibility of negative returns is not easily predicted from the harvest season price z-score. That is, years with high and low harvest season prices (relative to average) exhibit the phenomenon of negative returns. The calculated returns represent actual or *realized* returns, when in fact, the farmer's decision relies on *perceived* returns. If the farmer has experienced low or negative returns in the past in general, or low returns in years with high harvest prices, he may rely on those priors when choosing whether to store. Given household level data on farmer experience, and a minimum level of years of observations in each market, we could model farmer learning over time. The analysis we present here demonstrates that negative returns exist in all countries and time periods, and we are agnostic to heterogeneity in household risk preferences. Figure 1(c) shows that lower than average harvest prices (Quadrants I and IV) are associated with both negative and positive returns, and higher than average harvest prices (Quadrants II and III) are also associated with both negative and positive returns. A farmer could not easily foresee the expected returns to storage in a given year with information about the harvest season price alone. Figure 1(b) does show that higher mean seasonal returns are associated with markets with positive returns on average over time, therefore prior experience with frequent positive returns could influence farmer perception of returns to storage, but the possibility of negative returns remains a possibility in nearly all markets in our data.

2 Theoretical Model

To formally test for stochastic dominance, we need a simple model of household decision-making.

Household agricultural production is comprised of a single staple grain crop and the household can work off the farm. For simplicity, labor income is not included in this model. There are two time periods: period 1 is the end of the harvest season when the the household harvests the staple grain Q^*, which is the optimal production quantity from previously determined inputs. The household can sell the staple grain in the harvest season at price P_H, store it at a marginal cost of k with storage loss percentage $\alpha \in [0, 1]$, and consume c_H . Period 2 is the lean season when the household can consume or sell the staple grain if they chose to store in period 1, or purchase grain from the market for consumption. Lean season consumption is c_L. The household is a price taker in both input and output markets and complete markets exist for both, and storage S_H is restricted to being non-negative. At harvest time, the price of the staple grain is known P_H, but the lean season price is not known P_L.

Household income in the harvest and lean seasons, respectively:

$$Y_H = P_H\Big(Q^* - S_H - c_H\Big) - kS_H \qquad\qquad Y_L = P_L\Big((1-\alpha)S_H - c_L\Big)$$

The production quantity Q^* has been previously determined and c_H and c_L represent subsistence levels of grain consumption. All three parameters are exogenous to the decision to store. Gains (or losses) from stored grain are used to purchase the market good, whose price is normalized to 1. The lean season price is adjusted for inflation, under the assumption that the rate is exogenous.

The net payoff (i.e. change in wealth) from storing or not storing, if storing is restricted to be nonnegative:

$$w_{store} = -kS_H + P_L(1-\alpha)S_H \qquad\qquad w_{nostore} = P_H S_H$$

The household should store grain if $P_L(1-\alpha) - k > P_H$, however uncertainty about the lean season price can prevent the household from making the choice with the highest payoff. Thus the household is considering the trade off between storing and not storing for the purpose of agricultural incomes, and not explicitly for household consumption. Focusing on households that are net sellers allows us to avoid placing theoretical constraints on income and household preferences between grain and other goods.

While on average, the lean season price is higher than the harvest season price, the additional risk faced by the household may not be welfare-improving. Stochastic dominance (SD) allows us to determine whether the change in the distribution of payoff makes all risk averse households better off, and therefore justifies the decision to store. SD tests require limited assumptions on the household preferences, and have been used empirically to evaluate the welfare effects of policies and treatment effects. (Maasoumi and Heshmati, 2000; Millimet and Wang, 2011)

Assume the households have utility functions of the von Neumann-Morgenstern form. Let U_1 represent the set of utility function with $u' > 0$, i.e. monotonically increasing with respect to wealth, and U_2 represent the set of utility functions with $u'' < 0$, i.e. concave and risk averse. $F_1(w)$ and $F_2(w)$ represent the cumulative density functions (CDFs) under storing and not storing respectively, such that $F_{1,X}(w) = Pr(X < w)$ and $F_{2,Y}(w) = Pr(Y < w)$ for the random variables X,Y that represent the payoffs under storing and not storing.

Storing first order stochastically dominates (FOSD) not storing if and only if:

$$(F_2(y) - F_1(y)) \geq 0 \quad \forall y \in S$$

with strict inequality for at least one x. S represents the union of the supports of F_1 and F_2. If F_1 FOSD F_2, we can say that the expected welfare under storage is the same as or greater for all utility functions in U_1. FOSD implies SOSD and all higher orders, but the converse is not true.

Storing second order stochastically dominates (SOSD) not storing if and only if:

$$\int_{-\infty}^{y} (F_2(t) - F_1(t))\, dt \geq 0 \quad \forall y \in S$$

with strict inequality for at least one x. S represents the union of the supports of F_1 and F_2. If F_1 SOSD F_2, we can say that the expected welfare under storage is the same as or greater for all utility functions in U_2.

The behavior choice associated with uncertainty in a future period is prudence. Kimball (1990) defined prudence as the "propensity to prepare and forearm oneself in the face of uncertainty, in contrast to 'risk aversion' which is how much one dislikes uncertainty and would turn away from uncertainty if possible." The act of storing grain at harvest because of uncertainty about lean season grain prices is an example of prudence, or "precautionary demand for savings." Whereas risk aversion is associated with a concave vN-M utility function, positive precautionary savings requires a positive third derivative $u^{'''} > 0$. Thus all risk averse individuals are not prudent and quadratic utility is insufficient for precautionary savings. (Leland, 1968) Let U_3 represent the set of utility functions with $u''' > 0$.

Storing third order stochastically dominates (TOSD) not storing if and only if:

$$\int_{-\infty}^{y} \int_{-\infty}^{z} (F_2(t) - F_1(t))\, dt\ dz \geq 0 \quad \forall y \in S$$

with strict inequality for at least one x. S represents the union of the supports of F_1 and F_2. If F TOSD G, we can say that the expected welfare under storage is the same as or greater for all utility functions in U_3.

3 Estimation Strategy

3.1 Test of Equality

We first test for equality of the distributions $F_1(w)$ and $F_2(w)$ using a bootstrapped version of a Kolmogorov-Smirnov (K-S) statistic. McFadden (1989); Barrett and Donald (2003); Abadie (2002) The K-S test is a non-parametric approach to test the equivalence of two CDFs using empirical cumulative distribution functions derived from the data. The null hypothesis is that the two samples (payoffs under storing and not storing) come from the same underlying distribution. The sampling procedure used by Abadie allows for better identification of non-dominance in the tails of the distribution. We can reject the null hypothesis below a critical value (α). The full specification is included in Appendix A.1.

3.2 Stochastic Dominance

Second we use a generalized version of the K-S test to identify observed stochastic dominance rankings following Barrett and Donald (2003); Abadie (2002) and others. In order to test the significance of the observed rankings, we follow the suggested bootstrap approach of Linton et al. (2005); Maasoumi and Heshmati (2000) by resampling within each distribution separately. This approach provides a closer approximation of the true sampling distribution under the null hypothesis and allows us to estimate a confidence interval for the likelihood of the stochastic dominance event occurring. If the probability of dominance ranking exceeds a threshold of $1-\alpha$, we can be confident in the ranking. We include the full specification in Appendix A.2.

3.3 Parameters

Prices: The set of monthly maize prices for over 751 markets provides a rich dataset to characterize variation in the difference between planting and harvest season prices. The nominal monthly retail prices are converted to monthly real USD per kilogram. The harvest and lean season prices are described in Section 1.1. For example, if the primary harvest season is March-May and the planting season is October and November, the harvest season price (P_H) would be the real USD retail price in May and the lean season price (P_L) would be the higher of the real USD October or November retail prices for maize.

Storage costs: Positive returns to storage could be diminished by high storage costs. Many farmers use storage bags for which the cost is quite low. Indeed, the relatively low cost of storage technology is often cited as further evidence for the puzzle of farmers failing to take advantage of the inter-temporal arbitrage opportunity. The most common options are polypropylene (PP) bags which cost \$0.50USD and can store 50kg per bag or the Purdue Improved Crop Storage (PICS) bag which cost \$3.00 for 50kg capacity. (Walker et al., 2018) The estimated value for storage cost (k) is therefore \$0.10 USD per kilogram of maize stored. Burke et al. (2019) estimate storage costs at 2.5% of harvest price. We estimate storage cost at \$0.10 USD per kg of maize stored at harvest, consistent with the cost of the most basic storage method of PP bags.

Storage losses The probability of crop losses during storage varies with the crop, storage method, and stage at which storage is undertaken. Estimates of post-harvest losses (PHL) vary wildly. For purposes of the simulation, we consider storage losses as the quantity reduction in marketable maize after harvesting and before distribution: this covers losses while the grain is in storage due to pests, rodents or rotting. Kaminski et al. (2014) use self-reported data from six African countries and estimate post-harvest losses (PHL) of 1.4-5.9%. Gustavsson et al. (2011) report an average of 8% PHL for cereals in Sub-Saharan Africa. Burke et al. (2019) report relatively low estimates of 2.5% PHL for their study. Other studies using PP bags report storage losses in excess of 30%. (Hundie et al., 2019) We estimate storage losses (α)of 5.0% of total grain stored at harvest. This is a conservative estimate, as additional post harvest losses may occur during the distribution and marketing stage, further reducing farmers' realized returns.

4 Results: Stochastic Dominance

4.1 Test of Equality

Table 2 presents the results of the K-S test described in Section 3.1 for equality of the distributions. We can reject the null hypothesis of equal distributions for p-values less than our significance level of 0.1, which includes Burkina Faso, Burundi, Chad, Ethiopia, Malawi, Mali, Mozambique, Rwanda, Togo, Zambia and Zimbabwe. However, this test does not provide information on the level of inequality or direction of ranking between the distributions for storing and not storing.

4.2 Stochastic Dominance

We use empirical versions of the tests in Section 3.2 to identify and test for stochastic dominance. We use a Monte Carlo procedure that re-samples from the storage and non-storage payoff distributions separately to account for the small sample size and discrete data. Under this approach, if the probability of the event occurring is high, i.e. $\Pr(.) > 0.9$, and we observe stochastic dominance, this implies a 90% or higher confidence in the stochastic dominance results. The results are included in Table 2. Empirical CDFs are available from the authors upon request. For FOSD, there are observed rankings, but none of them are significant. For SOSD, storing grain is observed to dominant not storing in Burundi and Mali with more than 90% confidence. For TOSD, storing is also dominant in Zambia with more than 90% confidence.

For the remaining countries, we do not find significant evidence of a dominant strategy, despite observational rankings indicating stochastic dominance of both storing and not storing strategies at all levels. These values are conservative, given the low storage cost and loss parameters, and the selection of lean season prices as the maximum price a farmer would face during planting season. When the lean season price is calculated as the average planting season price, we do not obtain confidence in any level of dominance in any country.

Although there are observed instances of stochastic dominance at all orders, for both strategies, we only derive confidence in storing as a strategy in a limited number of countries, for risk averse farmers: Burundi and Mali, with storing dominant for prudent farmers in Zambia. Maize storage cannot stochastically dominate not storing, in any degree, in all countries in our data, and all times, due to the probability of negative returns to storage.

5 Discussion

Focus on average patterns of seasonal prices in the literature has led many researchers to overlook an important risk relevant to small farmer and small trader decision-making related to grain storage: years in which the lean season price fails to rise above the price at the time of harvest and no inter-temporal arbitrage opportunity occurs.

Our insight that storage implies a nontrivial probability of negative relative returns is consistent with farmers opting to sell immediately post-harvest if they have no hedging options. In the end, the result builds on explanations related to binding seasonal liquidity constraints, because those cause returns to lean-season sales to be relatively lower than returns to current sales. But the mechanism that we identify is different from previous analyses, arising from a different financial market failure. Our findings suggest the importance of analysis and interventions related to insurance and options associated with storage and lean season sales, both for small farmers and small traders in these markets. Forward contracts would protect against price risk, but are currently unavailable. If households could insure against price risk, might they store more?

Our focus on net sellers tells half the story: net buying household might have a greater tolerance for negative returns, in order to ensure sufficient food for the household in the lean season. But the result that storage is not a stochastically dominant strategy for all net sellers is a step towards changing the perception that credit alone is the solution to low storage uptake.

Price taking is a strong assumption, and allows us to avoid general equilibrium effects of storage choice on market prices. This is a first order approach focusing on the direct welfare effect of price fluctuations. Transaction costs would also add more complexity to the model, as they may induce farmers to opt in or out of the market, and vary seasonally based on regional transportation and trader networks. We demonstrate that using a conservative definition of negative returns and limited parameters, the assumption of positive returns to storage does not hold, and in fact, the possibility of negative returns might explain low storage uptake.

References

Abadie, A. (2002), 'Bootstrap Tests for Distributional Treatment Effects in Instrumental Variable Models', *Journal of the American Statistical Association* **97**(457), 284–292.

Aggarwal, S., Francis, E. and Robinson, J. (2018), 'Grain today, gain tomorrow: Evidence from a storage experiment with savings clubs in Kenya', *Journal of Development Economics* **134**, 1–15.

Barrett, C. B. (1996), 'On price risk and the inverse farm size-productivity relationship', *Journal of Development Economics* **51**, 193–215.

Barrett, C. B. (2007), 'Displaced distortions : Financial market failures and seemingly inefficient resource allocation in low-income rural communities', *Applied Economics* **i**(July), 1–13.

Barrett, G. F. and Donald, S. G. (2003), 'Consistent Tests for Stochastic Dominance', *Econometrica* **71**(1), 71–104.

Basu, K. and Wong, M. (2015), 'Evaluating seasonal food storage and credit programs in east Indonesia', *Journal of Development Economics* **115**, 200–216.

Burke, M., Bergquist, L. F. and Miguel, E. (2019), 'Sell low and buy high: Arbitrage and local price effects in Kenyan markets', *Quarterly Journal of Economics* **134**(2), 785–842.

Channa, H., Chen, A. Z., Pina, P., Ricker-Gilbert, J. and Stein, D. (2019), 'What drives smallholder farmers' willingness to pay for a new farm technology? Evidence from an experimental auction in Kenya', *Food Policy* **85**(March), 64–71.

Channa, H., Ricker-Gilbert, J., Shiferaw, F. and Abdoulaye, T. (2018), Helping Smallholder Farmers Make the Most of Maize through Loans and Storage Technology: Insights from a Randomized Control Trial in Tanzania.

Christian, P. and Dillon, B. (2018), 'Growing and Learning When Consumption Is Seasonal: Long-Term Evidence From Tanzania', *Demography* **55**(3), 1091–1118.

Finkelshtain, I. and Chalfant, J. A. (1991), 'Marketed Surplus under Risk: Do Peasants Agree with Sandmo?', *American Journal of Agricultural Economics* **73**(3), 557.

Gustavsson, J., Cederberg, C. and Sonesson, U. (2011), Food loss and food waste: Causes and solutions, Technical report, Food and Agriculture Organization of the United Nations.

Hundie, B., Abass, A. B., Hoeschle-zeledon, I., Mbwambo, H. and Bekunda, M. (2019), 'Exploring the profitability of improved storage technologies and their potential impacts on food security and income of smallholder farm households in Tanzania', *Journal of Stored Products Research* **82**, 98–109.

Kaminski, J., Christiaensen, L. and Gilbert, C. L. (2014), 'The End of Seasonality? New Insights from Sub-Saharan Africa', (February 2014), 1–44.

Kaminski, J., Christiaensen, L. and Gilbert, C. L. (2016), 'Seasonality in local food markets and consumption: Evidence from Tanzania', *Oxford Economic Papers* **68**(3), 736–757.

Kimball, M. S. (1990), 'Precautionary Saving in the Small and in the Large', *The Econometric Society* **58**(1), 53–73.

Leland, H. E. (1968), *Saving and Uncertainty: The Precautionary Demand for Savings*, Vol. 82, ACADEMIC PRESS, INC.

Linton, O., Maasoumi, E. and Whang, Y. J. (2005), 'Consistent testing for stochastic dominance under general sampling schemes', *Review of Economic Studies* **72**(3), 735–765.

Maasoumi, E. and Heshmati, A. (2000), 'Stochastic dominance amongst swedish income distributions', *Econometric Reviews* **19**(3), 287–320.

McFadden, D. (1989), *Testing for Stochastic Dominance*, fomby t.b. edn, New York, NY.

Millimet, D. L. and Wang, L. (2011), 'Is the quantity-quality trade-off a trade-off for all, none, or some?', *Economic Development and Cultural Change* **60**(1), 155–195.

Oi, W. Y. (1961), 'The Desirability of Price Instability Under Perfect Competition', *Econometrica* **29**(1), 58–64.

Oscar Maria Caccavale and Flämig, T. (2017), Collecting Prices for Food Security Programming, Technical report, World Food Programme (WFP), Rome.

Saha, A. and Stroud, J. (1994), 'A Household Model of On-Farm Storage under Price Risk', *American Journal of Agricultural Economics* **76**(3), 522.

Sahn, D. E. (1989), *Seasonal Variability in Third World Agriculture: The Consequences for Food Security*, Baltimore, MD (USA) Johns Hopkins Univ. Press.

Stephens, E. C. and Barrett, C. B. (2011), 'Incomplete Credit Markets and Commodity Marketing Behaviour', *Journal of Agricultural Economics* **62**(1), 1–24.

Stiglitz, J. E. (1969), 'Behavior Towards Risk with Many Commodities', *Econometrica* **37**(4), 660–667.

Walker, S., Jaime, R., Kagot, V. and Probst, C. (2018), 'Comparative effects of hermetic and traditional storage devices on maize grain : Mycotoxin development , insect infestation and grain quality', *Journal of Stored Products Research* **77**, 34–44.

Waugh, F. V. (1944), 'Does the Consumer Benefit from Price Instability?', *The Quarterly Journal of Economics* **58**(4), 602.

Table 1: Maize Retail Price Trends in Sub-Saharan Africa

Country	Years	Markets (1)	Market-Years (2)	Probability of Negative Returns (3)	Average Total Returns (4)	Average Positive Returns (5)	Average Negative Returns (6)
Benin	2007-2016	19	64	28.1%	14.1%	24.1%	-11.7%
Burkina Faso	2003-2019	19	226	19.5%	17.2%	22.5%	-4.5%
Burundi	2007-2019	64	246	29.3%	16.7%	30.4%	-16.5%
Cameroon	2005-2016	5	60	23.3%	16.3%	24.4%	-10.4%
Central African Republic	2004-2019	24	94	27.7%	28.9%	48.9%	-23.3%
Chad	2003-2018	12	77	15.6%	32.0%	40.6%	-14.7%
Cote d'Ivoire	2005-2019	10	61	29.5%	20.6%	35.4%	-14.8%
DR Congo	2008-2020	37	177	26.6%	25.8%	41.2%	-16.6%
Ethiopia	2006-2017	28	220	41.8%	20.4%	42.8%	-10.7%
Gambia	2006-2018	20	139	32.4%	27.4%	50.3%	-20.3%
Guinea	2017-2018	10	20	20.0%	42.9%	54.9%	-5.2%
Guinea-Bissau	2007-2018	2	12	50.0%	20.7%	57.0%	-15.7%
Kenya	2006-2019	9	80	38.8%	7.3%	17.2%	-8.2%
Malawi	2003-2019	71	762	19.7%	31.3%	41.8%	-11.6%
Mali	2003-2019	67	672	24.1%	29.8%	42.0%	-8.7%
Mozambique	2000-2019	24	325	10.5%	54.9%	62.6%	-11.0%
Niger	2000-2018	65	672	33.6%	12.7%	24.1%	-9.7%
Nigeria	2017-2019	4	9	55.6%	6.4%	28.6%	-11.3%
Rwanda	2008-2019	81	430	25.8%	20.2%	31.2%	-11.6%
Senegal	2007-2018	51	423	33.1%	12.9%	23.9%	-9.5%
Somalia	2009-2017	12	92	32.6%	15.7%	30.7%	-15.3%
South Sudan	2011-2019	9	45	40.0%	44.6%	97.2%	-34.3%
Togo	2001-2019	6	114	17.5%	27.9%	35.8%	-9.2%
Uganda	2011-2019	8	53	24.5%	31.2%	48.2%	-20.8%
Zambia	2003-2019	70	741	22.1%	32.2%	46.6%	-18.3%
Zimbabwe	2010-2018	24	108	21.3%	18.8%	26.8%	-10.9%
Total		751	5922	28.6%	24.2%	39.6%	-13.7%

[1] Monthly maize retail price data from the WFP Global Food Prices Database for 2000-2020 and agricultural season data from FAO-GIEWS.

[2] Prices were adjusted using IFS and FAOSTAT data on historical monthly exchange rates and CPI

[3] Columns (3)-(6): Returns are calculated for each "market-year" from the harvest season price to the subsequent lean season price. For markets with two agricultural seasons per year, returns for each year were averaged across the two seasons

Figure 1: Historical Trends for Maize Seasonal Returns

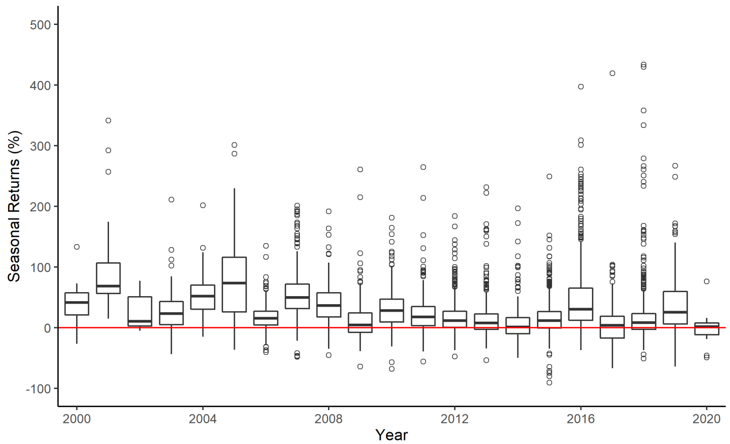

(a) Temporal Distribution of Negative Returns (market-year level)

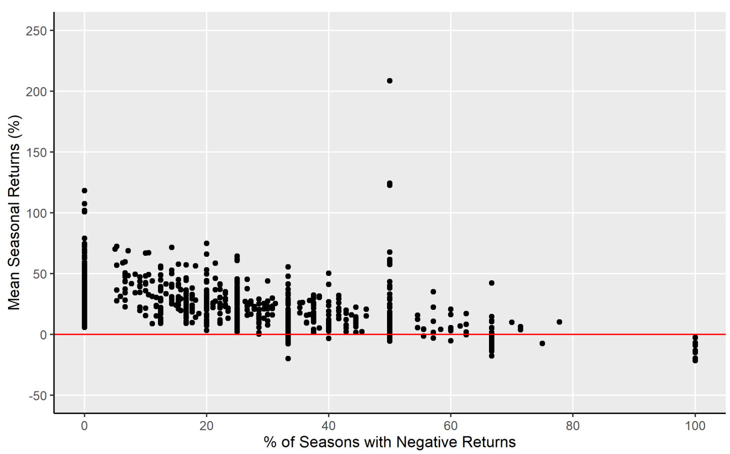

(b) Intensity and Frequency of Negative Returns (market-level)

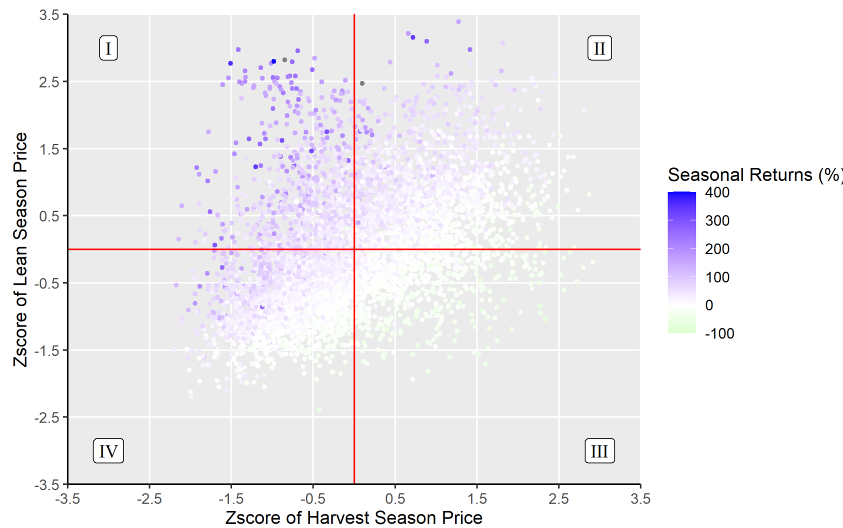

(c) Distribution of Negative Returns (market-year level)

Table 2: Stochastic Dominance Results

		First Order		Second Order		Third Order	
Country	K-S Test[1]	Rank[2]	$p(d^* \leq 0)^3$	Rank	$p(s^* \leq 0)$	Rank	$p(t^* \leq 0)$
Benin	0.28	no rank	0.09	storing	0.66	storing	0.70
Burkina Faso	0.05	no rank	0.11	storing	0.50	storing	0.58
Burundi	0.01	no rank	0.00	storing	0.97	storing	0.99
Cameroon	0.34	no rank	0.09	storing	0.66	storing	0.71
Central African Republic	0.11	no rank	0.01	no rank	0.31	no rank	0.36
Chad	0.00	no rank	0.26	storing	0.74	storing	0.76
Cote d'Ivoire	0.12	no rank	0.00	storing	0.73	storing	0.85
DR Congo	0.11	no rank	0.01	no rank	0.06	no rank	0.10
Ethiopia	0.07	no rank	0.00	storing	0.61	storing	0.66
Gambia	0.37	no rank	0.01	storing	0.55	storing	0.69
Guinea	0.52	no rank	0.22	no rank	0.35	no rank	0.42
Guinea-Bissau	0.98	no rank	0.12	no storage	0.54	no storage	0.63
Kenya	0.92	no rank	0.02	no rank	0.13	no rank	0.19
Malawi	0.00	no rank	0.33	no rank	0.28	no rank	0.30
Mali	0.00	storing	0.40	storing	0.94	storing	0.96
Mozambique	0.00	storing	0.68	storing	0.80	storing	0.81
Niger	0.41	no rank	0.02	storing	0.49	storing	0.56
Nigeria	0.96	no rank	0.25	no storage	0.62	no storage	0.72
Rwanda	0.01	no rank	0.00	storing	0.77	storing	0.80
Senegal	0.10	no rank	0.00	storing	0.79	storing	0.89
Somalia	0.98	no rank	0.02	no rank	0.29	no storage	0.37
South Sudan	0.30	no rank	0.01	no rank	0.28	no rank	0.40
Togo	0.07	storing	0.27	storing	0.74	storing	0.78
Uganda	0.18	no rank	0.04	storing	0.80	storing	0.84
Zambia	0.00	storing	0.33	storing	0.89	storing	0.93
Zimbabwe	0.05	no rank	0.04	storing	0.71	storing	0.80

[1] K-S Test: See Section 3.1. For $p < 0.1$, we can reject the null hypothesis of equality of the distributions

[2] Rank: See Section 3.2. If $\hat{x} \leq 0$, the distribution with $\hat{x} < 0$ is observed to be dominant, for $x = d, s, t$

[3] Probability: See Section 3.2. If $p(x^*) \leq 0$ exceeds 0.9, the distribution with $\hat{x} < 0$ is dominant, for $x = d, s, t$

A Stochastic Dominance Tests

A.1 Kolmogorov-Smirnov (K-S) Test Statistic

The two-sample K-S statistic is defined:

$$d_{eq} = \sqrt{\frac{n_1 n_2}{n_1 + n_2}} \sup_{y \in R} |F_{1,n_1}(y) - F_{2,n_2}(y)| \tag{A.1}$$

where n_1 and n_2 are the number of payoffs and F_1 and F_2 are the CDFs under storing and not storing, respectively.

The null hypothesis that the distributions under storing and not storing are equivalent:

$$F_{1,n_1}(y) - F_{2,n_2}(y) \quad \forall y \in R \tag{H.EQ}$$

We define empirical CDFs for F_1 and F_2 where $I(.)$ is an indicator function equal to 1 if the event occurs and 0 otherwise:

$$\hat{F}_{k,n_k}(y) = \frac{1}{N_k} \sum_{i=1}^{N_k} I(F_k \leq y) \quad \text{for} \quad k = 1,2 \tag{A.2}$$

We solve for $\hat{F}_{k,n_k}(y_j)$ where $y_j, j = 1...J$ are the points in N, the union of n_1 and n_2.

The test statistic:

$$\hat{d}_{eq} = \sqrt{\frac{n_1 n_2}{n_1 + n_2}} \max_j \left\{ \left| \hat{F}_1(y) - \hat{F}_2(y) \right| \right\} \tag{A.3}$$

Following Abadie (2002), we use a boostrapping procedure that resamples over the combined sample, splitting it into two subsamples with the same length as n_1 and n_2. For R reps, the p-value for the statistic is equal to:

p-value= $\frac{1}{R} \sum_{r=1}^{R} I(\hat{d}^*_{eq} > \hat{d}_{eq})$

We reject the null hypothesis of equality for p$<\alpha$ where $\alpha = 0.1$.

A.2 Observed Ranking and Hypothesis Testing for Stochastic Dominance

The K-S statistic can be modified for stochastic dominance, following McFadden (1989); Barrett and Donald (2003):

$$d = \sqrt{\frac{n_1 n_2}{n_1 + n_2}} \min \sup_{y \in Y} \{F_1(y) - F_2(y)\} \tag{B.1}$$

$$s = \sqrt{\frac{n_1 n_2}{n_1 + n_2}} \min \sup_{y \in Y} \int_{-\infty}^{y} (F_1(y) - F_2(y))\, dy \tag{B.2}$$

$$t = \sqrt{\frac{n_1 n_2}{n_1 + n_2}} \min \sup_{y \in Y} \int_{-\infty}^{z} \int_{-\infty}^{y} (F_1(y) - F_2(y))\, dydz \tag{B.3}$$

where the min is over $F_1(y) - F_2(y)$ and $F_2(y) - F_1(y)$.

The empirical CDFs can be determined as in EQ (A.2). The following are observed:

$$\hat{d} = \sqrt{\frac{n_1 n_2}{n_1 + n_2}} \quad \min\{maxd_1, maxd_2\} \tag{C.1}$$

where $y_j, j = 1...J$ are the points in N, the union of n_1 and n_2 and $d_1(y_j) = \hat{F}_{1,n_1}(y_j) - \hat{F}_{2,n_2}(y_j)$ and $d_2(y_j) = \hat{F}_{2,n_2}(y_j) - \hat{F}_{1,n_1}(y)_j$

$$\hat{s} = \sqrt{\frac{n_1 n_2}{n_1 + n_2}} \quad \min\{max(s_{1l}), max(s_{2l})\} \tag{C.2}$$

where $s_{kl} = \sum_{j=1}^{l} d_k(y_j)$ for k=1,2 and l=1..L.

and so forth for third order.

Storing is observed to FOSD not storing if $\hat{d} \leq 0$ and max $d_1 < 0$, and not storing FOSD storing if $\hat{d} \leq 0$ and max $d_2 < 0$. Similar interpreations are availabel for SOSD and $\hat{s}$ and TOSD and $\hat{t}$.

The null hypotheses for testing for stochastic dominance:

Storing first order dominates not storing if:

$$(F_2(y) - F_1(y)) \geq 0 \quad \forall \quad \forall y \in R \tag{H.1}$$

Storing second order dominates not storing if:

$$\int_{-\infty}^{y} (F_2(t) - F_1(t))\, dt \geq 0 \quad \forall y \in S \tag{H.2}$$

Storing third order dominates not storing if:

$$\int_{-\infty}^{y} \int_{-\infty}^{z} (F_2(t) - F_1(t))\, dt\ dz \geq 0 \quad \forall y \in S \tag{H.3}$$

In generating the test statistic, we follow a similar bootstrapping procedure as in the test of equality, but follow Maasoumi and Heshmati (2000); Linton et al. (2005); Millimet and Wang (2011) and others in resampling within n_1 and n_2 instead of the joint support, in order to avoid the 'null hypothesis bias'. The boundary between the null and alternative hypothesis is larger than the least favorable cases region (LFC) when $F_1 = F_2$ and therefore if it fails to hold while d,s,or t =0, the test may not have the correct asymptomic size and be biased. Resampling within n_1 and n_2 separately allows for approximation of the true sampling distribution. (Linton et al., 2005) We construct confidence intervals by evaluating whether the event $d* \leq 0$ has occurred for each iteration and then calculate the probability over all iterations. If $Pr(event)$ is greater than $1 - \alpha$ where $\alpha = 0.1$, and the event was also observed, we can be confidence that first, second, or third order dominance has occurred.