AN EXPERIMENTAL APPROACH TO TESTING THE COMPETITIVE STORAGE MODEL

SIMONE PFUDERER

Contributed Paper prepared for presentation at the 89th Annual Conference of the Agricultural Economics Society, University of Warwick, England
13 - 15 April 2015

Abstract

A new approach to testing the competitive storage model is introduced. A relatively simple model is taken to the experimental laboratory. Market outcomes in the experiment deviate from the predictions of the competitive storage model in a number of ways. Average storage and the variability of storage are below the levels predicted by the competitive storage model. The resulting price series, therefore, tend to be more variable than would be the case if stockholders behaved according to the competitive storage model. In addition, the predicted relationship between availability and storage is non-linear but is linear in the experiment.

Keywords: commodity markets, prices, storage, experiment

JEL classification: C61, C92, Q11

1 Introduction

It is undeniable that stockholding affects prices. The role of low stocks in recent price movements, especially the 2007/08 and 2010/11 grain price peaks, has been widely discussed. Global stocks-to-use ratios were relatively low in the mid-2000s for a number of grains and oilseeds. Several studies, therefore, suggest that low stocks contributed to price peak of 2007/08 (e.g. European Commission, 2008; OECD, 2008; Wiggins and Keats, 2010).

The important role of storage in price formation is well known and a number of models have been developed over time. Gustafson (1958) introduced supply shocks in form of probabilistic harvests into a model of intertemporal price formation and storage. The storage decision is made taking into account the amount of the good available in the current period and the expected supply in the following period. The distribution of the harvest is known but the outcome of the stochastic harvest is only known for periods up to the current period. The model is solved for the profit-maximising storage using numerical methods. Gustafson's model is basis of the modern stochastic competitive storage models which remain the main tool in theoretically analysing the stocks-price relationship.

This paper introduces a new approach to testing the competitive storage model. A relatively simple version of the competitive storage model is taken to the experimental laboratory. The advantage of this approach is that the characteristics of the model are under the control of the experimenter and behaviour can be studied in this controlled environment. In the real world, information on storage decisions is difficult to collect. In an experimental setting, many parameters that have to be estimated in the real world and which might be only known with error can be controlled. In addition, data collection is straightforward in the laboratory.

There is little pre-existing experimental literature on storage. The only instance of which I am aware is Abbink et al. (2011) who present an experimental study of storage decisions of maize traders in Zambia. The model on which their experimental design is based is, therefore, a game of strategic interaction between government and a small number of private stockholders.

In this experiment participants were asked to make storage decisions within a competitive storage model framework. Their behaviour, the aggregate stock level and the resulting price series are compared to the model assumptions and predictions.

2 The competitive storage model

The competitive storage model is a rational expectations model of optimal storage for a commodity where production is uncertain and the commodity is storable from one period to another, such as an annual agricultural crop that is storable. The model can be formulated with a single state variable, a single control variable and one arbitrage equation.

In every period t, stockholders start with a pre-determined level of the commodity that was carried forward from the previous period, s_{t-1}. The quantity of the commodity available in period t, availability a_t, is the state variable. Availability is the quantity carried forward from the previous period, s_{t-1}, plus the harvest in period t. The quantity harvested is an exogenous random variable ϵ_t. Availability in period t, a_t, can be used for consumption and storage. The amount used for consumption, c_t, is sold to consumers at the price that clears the market, $p_t = P(c_t)$. The difference between availability, a_t, and consumption, c_t, is the amount of storage s_t, the control variable in the model. In the next period t+1, the pre-determined stock level carried forward from the previous period is s_t, which together with exogenous production, ϵ_{t+1}, gives the total amount available in period t+1, a_{t+1}. The transition equation of the dynamic model is therefore:

$$a_{t+1} = s_t + \epsilon_{t+1} \tag{1}$$

Availability in t+1, a_{t+1} depends on the exogenous variable ϵ_{t+1} and the endogenous variable s_t.

The stockholders' objective is to maximise their expected profit which leads to the storage arbitrage equation.

$$\mathrm{E_t}[\mathrm{p_{t+1}}] - p_\mathrm{t} = \pi_t \tag{2}$$

where $E_t[p_{t+1}]$ is the expected price in period t for period t+1, p_t is the price in period t and π_t is expected marginal profit of a unit stocks. As noted before, s_t the amount of the commodity stored from period t to period t+1, is the control variable that the agents in the model, the stockholders, adjust to maximise their profits. In a competitive market and in the absence of storage capacity limitations, stockholders will adjust the storage level until the expected profits are zero. The arbitrage equation can be written as a function of the state variable a and the control variable s.

$$\mathrm{E_t}[P(a_{t+1} - s_{t+1})] - \mathrm{P}(a_t - s_t) = \pi_t \tag{3}$$

As a whole, the economy cannot borrow production from future periods. Therefore, storage cannot be negative. The non-negativity constraint limits arbitrage when storage is zero and in these situations expected profits are negative.

$$\mathrm{E}_t[P(a_{t+1} - s_{t+1})] - \mathrm{P}(a_t - s_t) = \pi_t \tag{4}$$

$$s_t \geq 0 \Rightarrow \pi_t \leq 0, \; s_t > 0 \Rightarrow \pi_t = 0$$

When expected profits are positive, that is when the expected price in period t for period t+1 exceeds the current price, stockholders will store another unit which will lower availability and increase the price in period t and increase availability and decrease the price in period t+1. Stockholders will continue to increase storage until the point is reached where expected profits are zero. When profits are negative, that is when the expected price in period t for period t+1 is lower than the current price, stockholders will reduce storage. Reducing storage increases availability and reduces the price in period t and, at the same time, lowers availability and increases the price in period t+1. When stocks are zero, arbitrage is limited and expected profit from storage is negative.

With a maximum storage capacity in the economy, the possibility of increasing storage when expected profits are positive is limited by the maximum storage level. When storage is at the maximum storage level, therefore, expected profits are positive.

$$\mathrm{E}_t[P(a_{t+1} - s_{t+1})] - \mathrm{P}(a_t - s_t) = \pi_t \tag{5}$$

$$s_t > 0 \Rightarrow \pi_t \geq 0, \; s_t < s_{max} \Rightarrow \pi_t \leq 0, 0 < s_t < s_{max} \Rightarrow \pi_t = 0$$

In a rational expectations model, agents' expectations for variables in the next period have to be consistent with the resultant distribution of these variables given the structure of the model, the parameters in the model and the expectations. In the present model the expectations are with respect to the price in the next period. In this simple version of the model, stockholders are the only agents. They have to make a decision on how much of the commodity to store from one period to the next and this decision depends on their expectations of the price in the next period. The model does not have a closed-form solution because of the non-negativity constraint on storage and needs to be solved numerically.

3 The experiment

The experiment was run at the Cognitive and Experimental Economics Laboratory (CEEL) of the University of Trento in April 2014. Overall 88 participants took part in the experiment which formed 11 groups of eight, leading to 11 independent market observations. The experiment was run as a computerised experiment, programmed and conducted using the z-Tree software (Fischbacher, 2007).

At the start of the experiment, participants were randomly assigned to groups of eight. The composition of the group did not change during the experiment ("partner matching"). Each session consisted of 25 periods which falls within the range of the number of rounds used in asset market experiments following Smith et al. (1988) that generally consists of 15 to 30 periods (Noussair and Tucker, 2013). Given possible learning and end game effects, a number of periods towards the higher end of this range was chosen. In a trial of the experiment, feedback from the volunteers suggested that 25 rounds was an adequate length for the experiment, some even suggesting that it could be longer. With 25 rounds 275 observations at the group level and 2200 at the individual level were collected.

At the beginning of the experiment, each participant received an endowment of 10 Experimental Currency Units (ECU). All participants in the experiments had exactly the same role, namely that of wheat traders, who can buy, sell and store wheat. The eight participants in each group participated in the same market. At the start of the experiment stocks were zero for all participants. It would have been possible to randomly assign units of wheat to participants at the start. However, this would have introduced another random process to the experiment which would have had to be explained to participants creating an unnecessary additional level of complexity and possibly leading to confusion of the random processes with adding little to the experiment.

Each participant had a capacity of storage of one unit leading to a maximum storage capacity at the market level of eight. Wheat could only be bought, sold and stored in full units. In each period the participants were therefore either potential buyers (those participants that did not carry forward a unit of stock from the previous period) or potential sellers (those participants that carried forward one unit of stock from the previous period). As a consequence, each participant had to make a straightforward decision in each round, namely, for potential sellers, from which price on to sell and for buyers up to which price to buy.

There was a single wheat harvest in each round which followed a simple three point distribution with a small harvest of 50 units, a medium of 60 units and good harvest of 65 units. In each round, the probability that the harvest was 50 units was 20 per cent, that it was 60 units was 40 per cent and that it was 65 units was 40 per cent. All units available in a round were used, either for consumption or for storage to the next round.

Those participants who had carried forward a unit of wheat from the previous round were potential sellers in that round and were asked to submit the minimum price for which they wanted to sell their unit of wheat (in steps of 0.05 ECU). The participants who had not carried forward a unit of wheat from the previous round were potential buyers in that round and were asked to submit the maximum price for which they wanted to buy a unit of wheat (in steps of 0.05 ECU).

Consumption depended on the price and was specified in the deterministic consumption-price function, which was communicated to the participants.

The consumption function in the experiment was:

$$C(p_t) = 90 - 20 * p_t \qquad (6)$$

where C is consumption and p_t is the price in period t.

An algorithm included in the z-Tree program established the market price.

4 Competitive storage model predictions

In this section, model predictions are derived using a discrete dynamic model with the same parameters as in the experiment.

Figure 1 shows the storage level at each possible level of availability according to the competitive storage model prediction. The storage function requires that up to a level of availability (harvest plus stocks carried forward) of 58 storage will be zero and that from a level of availability of 70 storage will be at its maximum level of eight units.

Figure 1 Equilibrium storage function

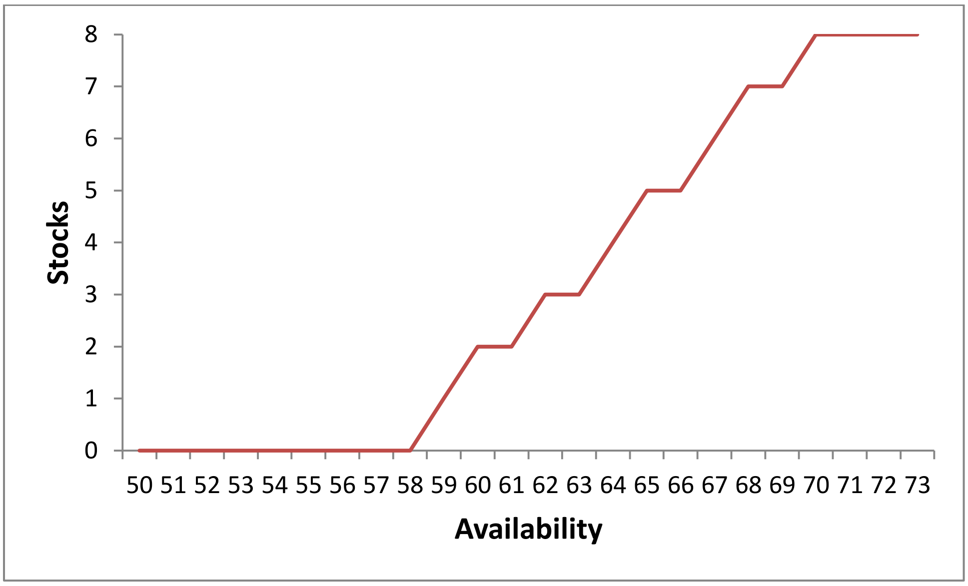

Figure 2 shows the equilibrium price function (solid line) and the price function without storage (dashed line). It shows that up to the level of availability of 58 when no storage takes place, the price is the same with and without storage. When availability exceeds 58 units, the optimal storage rule leads to storage, leaving less for consumption, and thus leads to an increase in the price compared to the no storage scenario. From 70 units of availability onwards the equilibrium price and price without storage lines are parallel. At 70 units of availability optimal storage reaches the maximum storage capacity level of eight units and storage cannot increase further at higher availability levels.

Figure 2 Equilibrium price function and price function without storage

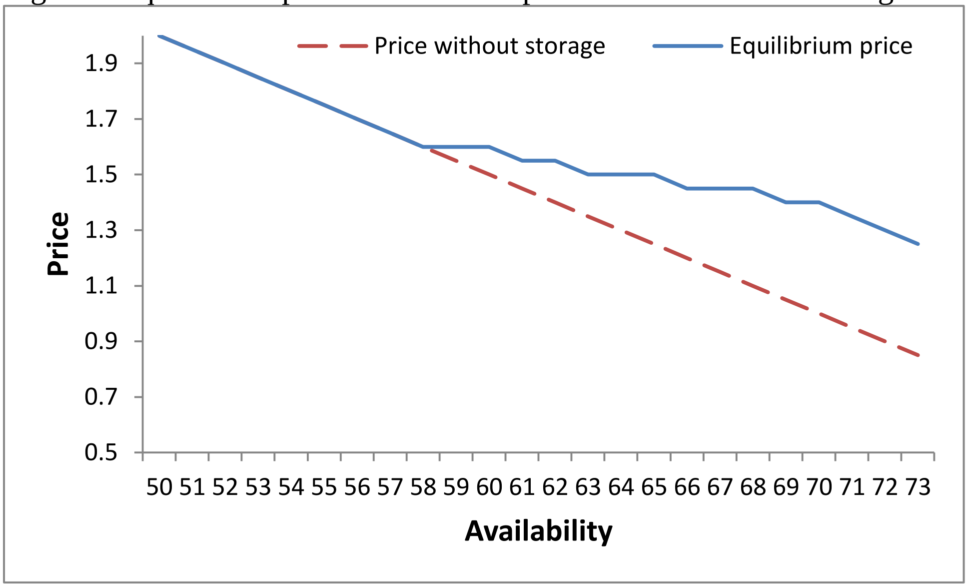

The optimal storage[1] outcomes, to which the experimental results are compared, are based on simulations of the experiment assuming that participants behave as close to the competitive storage model outcomes as possible within the experimental design. The approach taken is explained in the following.

For each level of availability the model solution gives the optimal storage, consumption, price and expected price. Table 1 shows the availability levels possible in the experiment, the corresponding stock levels and expected prices at these levels of availability according to the competitive storage model solution.

Table 1 Competitive storage model: availability, storage and expected price

Availability	50	51	52	53	54	55	56	57	58	59	60	61
Storage	0	0	0	0	0	0	0	0	0	1	2	2
$E_t(P_{t+1})$	1.64	1.64	1.64	1.64	1.64	1.64	1.64	1.64	1.64	1.59	1.58	1.58
Availability	62	63	64	65	66	67	68	69	70	71	72	73
Storage	3	3	4	5	5	6	7	7	8	8	8	8
$E_t(P_{t+1})$	1.55	1.55	1.52	1.51	1.51	1.46	1.43	1.43	1.40	1.40	1.40	1.40

As a consequence of the price function and the limitation of storage to full units, the price can only take values in steps of 0.05 ECU whilst the expected price can take values in between those possible in the experiment. If all participants follow the optimal strategy, participants maximise their profit if they submit the expected price. However, participants had to submit limit prices that are possible in the experiment, i.e. prices in steps of 0.05 ECU. In the simulation of optimal storage in the experimental setting, therefore, for buyers not to make an expected loss they submit a rounded-down expected price and sellers a rounded-up expected price. Table 2 shows the optimal decision rule for buyers and sellers in the experiment to which actual behaviour will be compared. This strategy maximises earnings within the experiment.

Table 2 Approximation of optimal limit prices for buyers and sellers

Availability	50	51	52	53	54	55	56	57	58	59	60	61
Submitted price buyer	1.60	1.60	1.60	1.60	1.60	1.60	1.60	1.60	1.60	1.55	1.55	1.55
Submitted price seller	1.65	1.65	1.65	1.65	1.65	1.65	1.65	1.65	1.65	1.60	1.60	1.60
Availability	62	63	64	65	66	67	68	69	70	71	72	73
Submitted price buyer	1.55	1.55	1.50	1.50	1.50	1.45	1.40	1.40	1.40	1.40	1.40	1.40
Submitted price seller	1.55	1.55	1.55	1.55	1.55	1.50	1.45	1.45	1.40	1.40	1.40	1.40

[1] In the following whenever reference is made to optimal strategy or behaviour, reference is made to these simulation results.

5 Results

Figure 3 shows the random harvest outcome of the four sessions. In each session, the harvest outcome was the same for all groups. Before starting the analysis, checks were carried out for a possible framing effect of the two different versions of the instructions. In both versions the same examples were used but a different example was described in detail in the two versions. Sessions 1 and 3 described an example with a lower price submitted and sessions 2 and 4 with a higher price submitted. Sessions 3 and 4 started with the same level of the harvest – namely 60 – and thus the prices submitted can be directly compared. The average of the prices submitted in the first round of session 3 was 1.71 ECU and that of session 4 was 1.78 ECU, both higher than 1.50 ECU, the price used in the example that was a possible market price outcome with within the experiment (the other price in the example was not a possible market price but could be submitted as minimum or maximum price). The difference in the mean price of session 3 and session 4 in the first round is not statistically significant.

Figure 3 Harvest outcome series of the four sessions

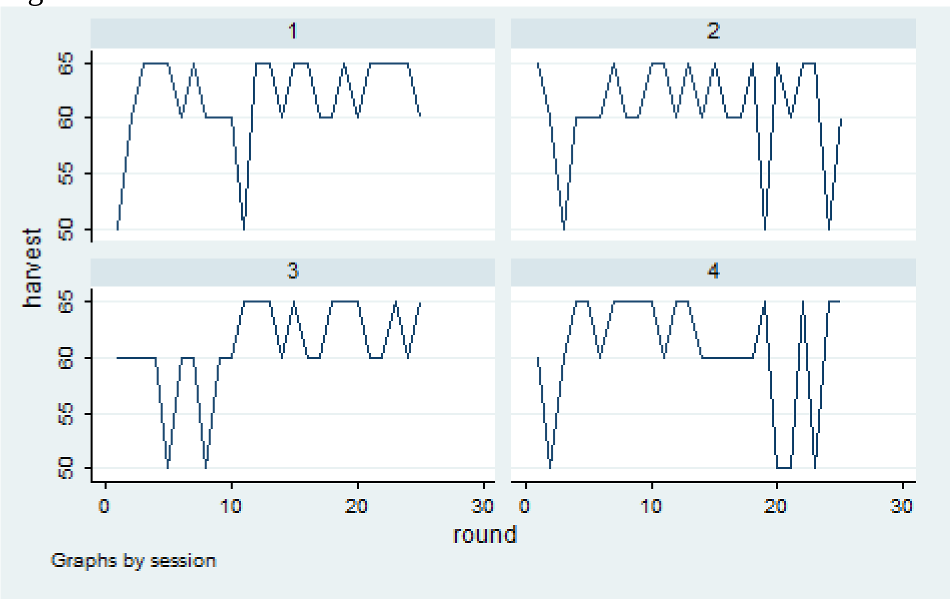

A comparison of the submitted price across all four sessions is not meaningful because the submitted prices will be influenced by the level of the harvest outcomes. A comparison of the deviation from the optimal price shows that in all sessions the average submitted prices were higher than the optimal prices. The deviation is greater for sessions 2 and 4 than for Sessions 1 and 3 but not statistically significantly so. Looking at the first two periods, the deviation is greater for Sessions 1 and 3 than for sessions 2 and 4 but not significantly so.

These results suggest that there is no statistically significant effect from the examples used in the instructions.

In the experiment, participants with a unit of wheat in storage at the end of round 25 were paid the average price for this unit. Therefore, in round 25 the expected price for the next period is the average price over the last 24 rounds. This average price is close to, but not exactly the same as, the expected price assumed in the simulations. Therefore the results of the last round are not directly comparable to those of previous rounds but the effect should be small. There might also be learning effects especially at the start of the experiment. However, the main results are robust to the exclusion of the first three and last two periods. The results presented in the following are based on all 25 rounds but they do not substantially differ when the first three and last two periods are excluded.

5.1 Comparison of actual and optimal storage

Figure 4 shows the optimal storage level and the actual storage levels by the groups in each session.

Figure 4 Optimal and actual storage level by session

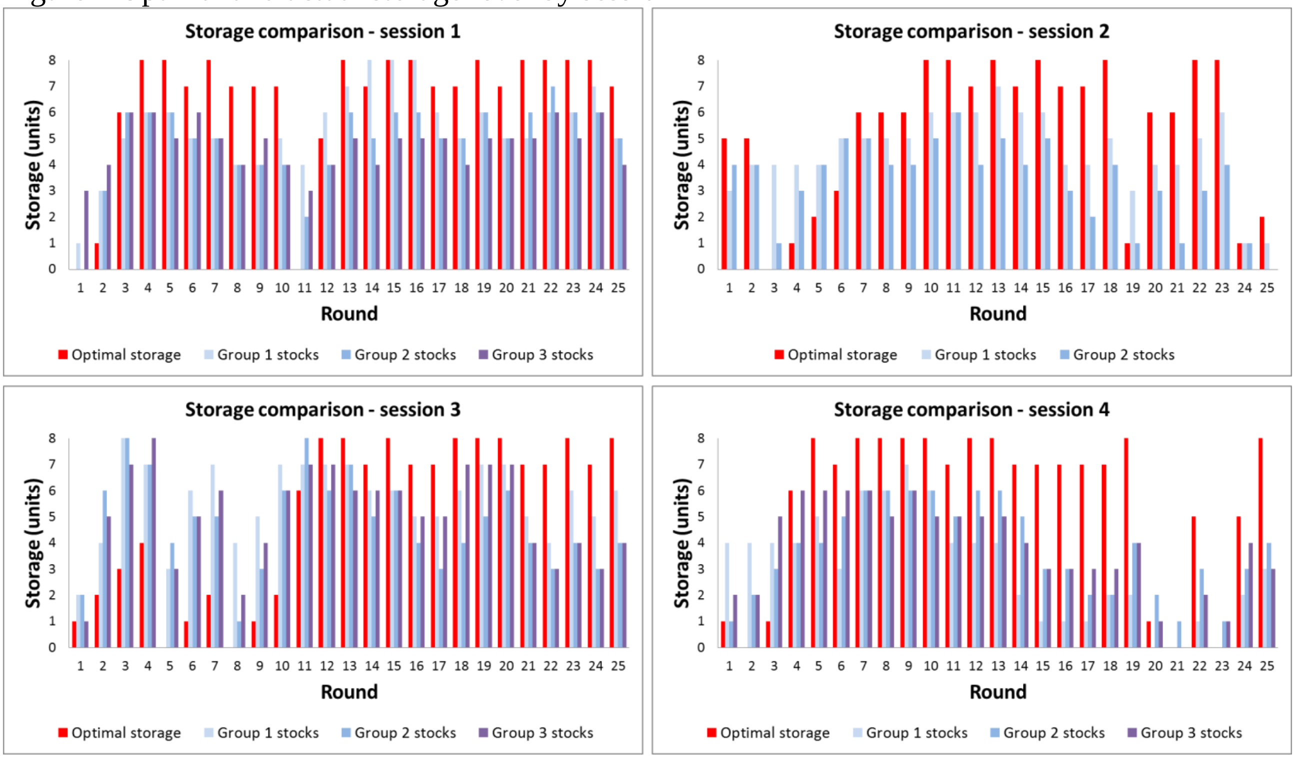

It is clear from Figure 4 that in general storage is below its optimal level when the optimal level is high but that it is above the optimal level when optimal storage is low. Table 3 compares the means and standard deviations of storage for the four sessions with the optimal mean and standard deviation. The optimal mean and standard deviations are based on the approach explained above.

In sessions 1 and 4 the mean stock levels for all groups are statistically significantly below the optimal level at the 5% significance level. In session 2 for group 2 the equality of mean is rejected at the 1% level but for group 1 the hypothesis that the mean is equal to the optimal mean cannot be rejected. In session 3, the hypothesis that the mean stock level is equal to the optimal mean stock level cannot be rejected for any of the groups. Excluding the first three (possible learning effects and effect of starting with zero storage) and the last two periods (possible end game effect) does not significantly change these results.

Table 3 Optimal and actual mean and standard deviation of storage (rounds 1 to 25)

	Optimal	Group 1	Group 2	Group 3
Session 1				
Mean	6.52	5.44	4.92	4.76
p-value		0.0357**	0.0042***	0.0010***
Std. Deviation	2.45	1.58	1.53	0.88
p-value		0.0184**	0.0119**	0.0000***
Session 2				
Mean	5.36	4.52	3.40	
p-value		0.0906*	0.0017***	
Std. Deviation	2.71	1.48	1.58	
p-value		0.0021***	0.0054***	
Session 3				
Mean	5.12	5.68	4.76	5.12
p-value		0.7918	0.3077	0.5000
Std. Deviation	3.07	1.46	1.79	1.81
p-value		0.0003***	0.0051***	0.0060***
Session 4				
Mean	5.60	3.04	3.72	3.80
p-value		0.0006***	0.0055***	0.0055***
Std. Deviation	3.06	2.05	1.74	1.78
p-value		0.0282**	0.0040***	0.0052***

Notes: The p-values are for the one-sided t-test that the actual mean is equal against the hypothesis that it is lower than the optimal mean and the one-sided F-test that the actual standard deviation is equal against the one-sided hypothesis that it is smaller than the optimal standard deviation. *** denote rejection at the 1% level, ** at the 5 % level.

The hypothesis that the actual standard deviation of storage is equal to the optimal standard deviation of storage is rejected at the 5 % level for all groups, including those in session 3 for which the hypothesis that the means are equal could not be rejected. Thus, the results suggest that in general stock levels are lower than would be optimal and that in all cases stock levels vary less than would be optimal. The next section looks at the impact of the storage behaviour on the price series.

5.2 Comparison of standard deviation of actual, optimal and no-storage price series

While for storage a comparison of actual stocks to a no storage scenario is not meaningful, for price series this comparison provides interesting insights into whether or not storage significantly reduces the variation of prices in the experiment. The expectation is that the actual price series in the experiment would lie somewhere between the price series that one would get without storage and the price series one would get with optimal storage. Because of the linear price function, the mean price will be very similar for optimal storage, no-storage and for the actual price series. The only difference in the mean price over the 25 rounds is due to the fact that storage at the end of the 25 rounds differs. If all units in storage were to be sold in the last round, the mean price would be exactly the same. The focus of the analysis is therefore on the standard deviation of the price series. Figure 5 shows the optimal, no-storage and actual price series of the four sessions.

Figure 5 Optimal, no-storage and actual price series by session

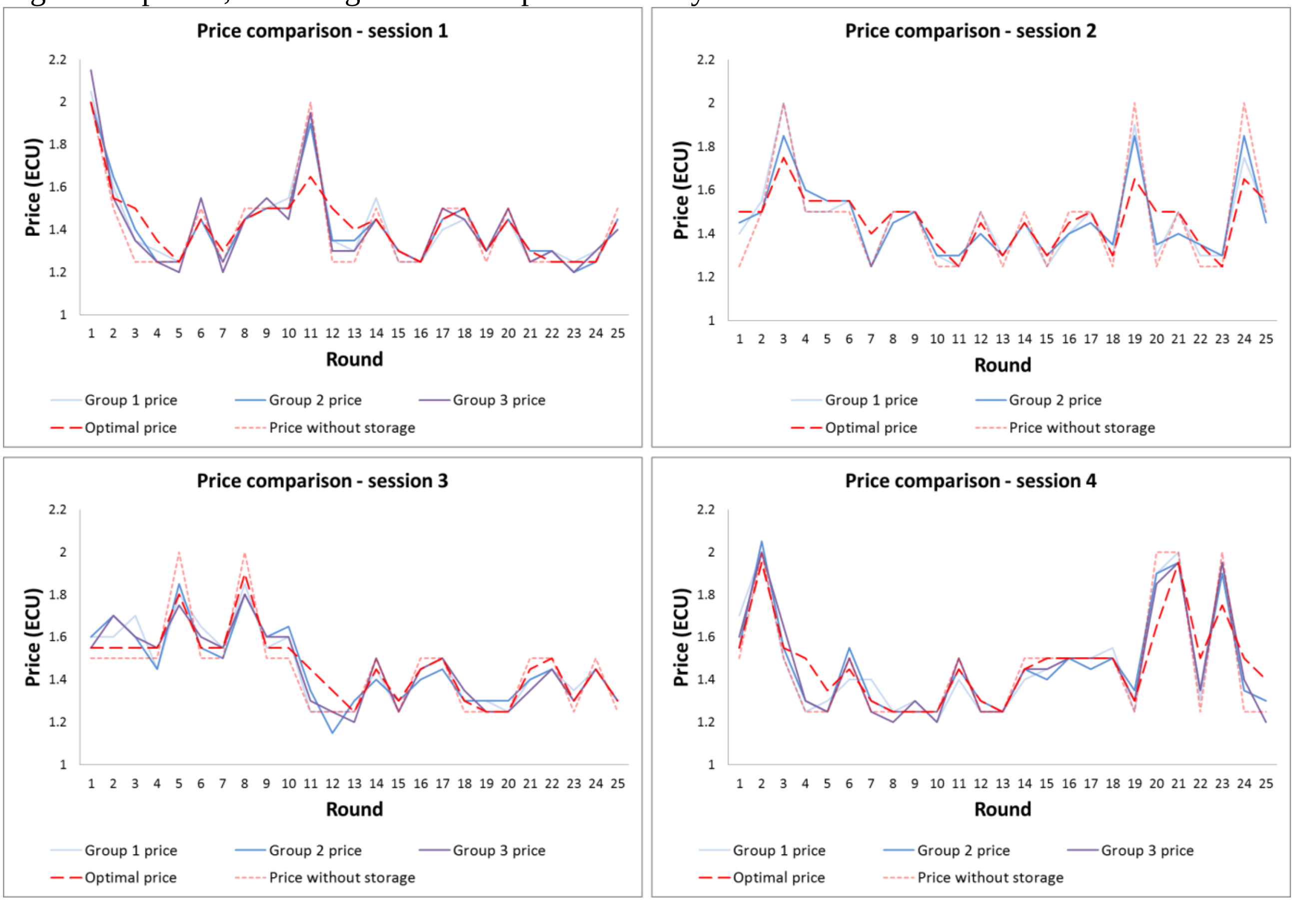

Figure 5 suggests that the optimal price series differs especially when prices are high. Table 4 reports the standard deviation of the price and the test results for the null hypothesis that the standard deviations of the actual price series are equal to those of the optimal price series against the one-sided hypothesis that the actual standard deviation is higher than the standard deviation of the optimal price series. The tests comparing the standard deviation of the actual price series with the no-storage price series test the hypothesis that the standard deviations are equal against the hypothesis that the actual standard deviation is lower than the standard deviation without storage.

As expected, the actual standard deviation of the price is above the standard deviation of the price for the optimal strategy for all groups. However, for one group (session 1, group 3), the standard deviation of the price is also higher than that of the no-storage price. Thus, the storage behaviour of this group made the price more variable than would have been the case without any storage.

Table 4 shows that the only occurrence when the standard deviation of the actual series is statistically significantly below the series of no-storage (at the 10% level) is when it is also statistically significantly above the standard deviation of the optimal storage price series (at the 10% level). For three of the eleven groups the standard deviation is statistically significantly above the optimal price series. It is interesting that for session 3 the standard deviations of the price series without storage and that with optimal storage do not differ statistically significantly. It is therefore not surprising that no statistically significant results were found for the standard deviation of the actual price series as the latter is expected to lie between the optimal and no-storage benchmarks.

To conclude, in the experiment storage does not achieve the optimal reduction in the standard deviation that would be possible if participants behaved according to the competitive storage model. In ten out of the eleven groups, the standard deviation of the price series is not statistically significantly below the standard deviation of the price series that would occur without storage.

Table 4 Optimal, no storage and actual standard deviation of the price series (rounds 1 to 25)

	Optimal	No storage	Group 1	Group 2	Group 3
Session 1					
Standard deviation	0.1646	0.2151	0.2051	0.1942	0.2240
p-value optimal		0.0990*	0.1441	0.2124	0.0692*
p-value no storage			0.4095	0.3105	0.5784
Session 2					
Standard deviation	0.1279	0.2358	0.1927	0.1726	
p-value optimal		0.0020***	0.0249**	0.0746*	
p-value no storage			0.1646	0.0666*	
Session 3					
Standard deviation	0.1612	0.2041	0.1769	0.1742	0.1718
p-value optimal		0.1272	0.3258	0.3533	0.3787
p-value no storage			0.2446	0.2217	0.2022
Session 4					
Standard deviation	0.1926	0.2669	0.2503	0.2384	0.2465
p-value optimal		0.0585*	0.1034	0.1515	0.1170

p-value no storage	0.3776	0.2923	0.3500

Notes: The p-values are for the one-sided F-test that the actual standard deviation is equal against the hypothesis that it is higher than the optimal standard deviation and the one-sided F-test that the actual standard deviation is equal against the one-sided hypothesis that it is lower than the no storage standard deviation. *** denote rejection at the 1% level, ** at the 5 % level and * at the 10% level.

5.3 Decisions at the group level

The decisions by participants in each round and the random harvest outcome determine the stock level at the end of the round and the price in the round. In the competitive storage model, the stock level at the end of the round and the price in the round only depend on availability. Figure 6 plots the price against availability for the optimal strategy and for all sessions in the experiment. Whilst the relationship between availability and price in the experiment looks fairly linear, this is not the case for the optimal availability-price relationship which has two kinks, one at the level of availability where storing one unit becomes optimal and one at the level of availability where the maximum storage level becomes optimal. In between these two kinks the slope is much flatter.

Figure 6 Optimal and actual availability-price relationship

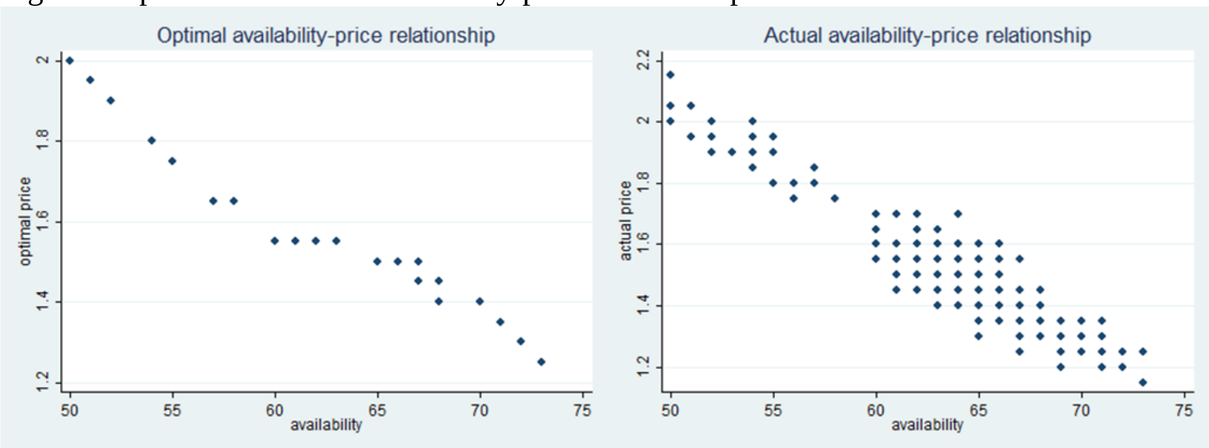

Figure 7 plots stocks at the end of the round against availability for the optimal behaviour and for the experiment. While optimal storage behaviour leads to left and right censoring of this relationship, this feature is not apparent in the case of the actual availability and storage relationship.

Figure 7 Optimal and actual availability-stock relationship

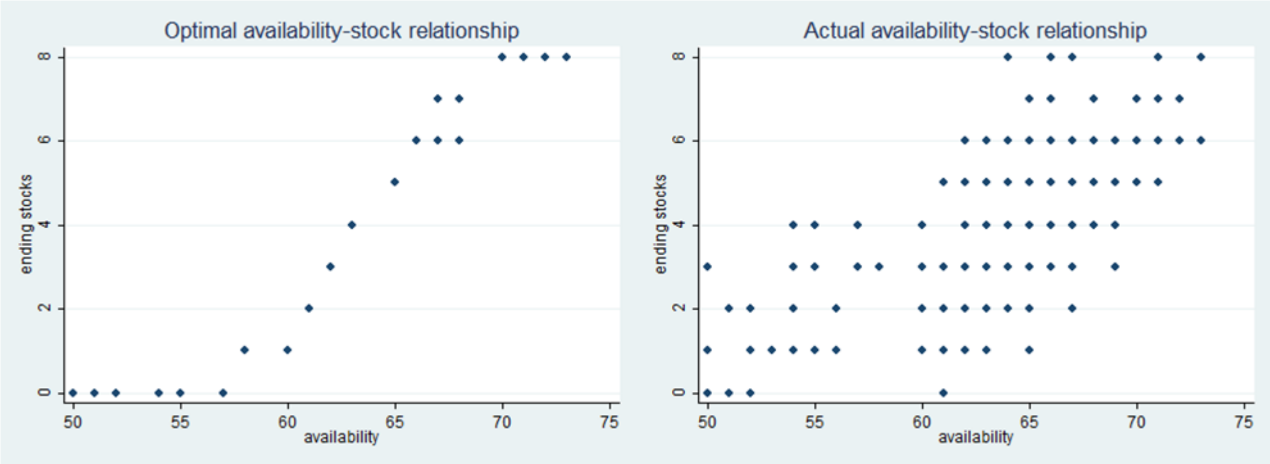

Similarly as for the availability-price relationship, while the optimal availability-storage relationship is clearly not linear, this is not obvious for the actual availability-storage relationship. Whether or not a linear model fits the availability-storage and the availability-price relationship can be tested using a Ramsey Regression Equation Specification Error Test (RESET) test (Ramsey, 1969).

Before analysing the availability-storage decision in more detail, the implications of the experimental design and the optimal strategy within the experimental design need to be investigated. Making the optimal strategy dependent on whether the participant is a buyer or a seller in that round has important implications for the availability-storage relationship. In the competitive storage model carry-in stocks and harvest have the same impact on storage. For the optimal strategy in the experiment, this is not the case, because the price submitted depends on the role of the participant.

If we have an availability of say 65, the expected price from the model solution is 1.51 ECU. The availability of 65 can either be the result of a harvest of 65 and no stocks carried forward or a harvest of 60 and 5 units of stocks carried forward. In the first case there are only buyers and in the second there are 5 sellers and 3 buyers. In the model these two scenarios are equivalent. However, for the optimal strategy within the experiment, the two scenarios are not equivalent because buyers submit a rounded-down optimal price and seller a rounded-up optimal price. In the first case all participants submit a price of 1.50 ECU but in the latter case five submit a price of 1.55 ECU and three a price of 1.50 ECU. So in the latter case stocks will tend to be higher for the same level of availability.

If sellers and buyers were both to round up or both to round down, the additional dependence on stocks would disappear and the impact of carry-in stocks and harvest would be identical. This would be in some ways preferable but it is not assumed in the analysis here because within the experimental design the suggested optimal strategy is superior in terms of expected profit. If sellers and buyers were both to round down, sellers would make a negative expected profit, which they could avoid by rounding up. If sellers and buyers were both to round up, buyers would make a negative expected profit, which they could avoid by rounding down. As a consequence of this assumption, for the optimal strategy proposed, the impact of one unit of wheat from the harvest and one unit of the carry-in stocks is not the same.

For this reason, the RESET test is based on a linear fixed effect regression[2] including both harvest and carry-in stocks and not just their sum, i.e. availability. The regression model of the test also includes squares and cubes of the fitted values to test if the linear model fits the data. The F-statistic testing that the coefficients of squares and cubes of the fitted values are jointly zero is reported together with the corresponding p-value. The test results are reported for the regression of price on carry-in stocks and harvest and that of storage on carry-in stocks and harvest. The test is carried out for all sessions together and for each session individually.

Table 5 shows the F-statistic and p-values for the RESET test. These results confirm that the relationships are clearly non-linear for the optimal strategy but linear in the experiment. For the actual relationships a linear model seems the appropriate specification. The only case where the RESET test does not reject the hypothesis that the squares and cubes of the fitted values are jointly zero is for the regression on price in session 3. Here the null hypothesis is rejected at the 10% level.

Table 5 RESET test results

	All sessions	Session 1	Session 2	Session 3	Session 4	Optimal
Dependent variable – price						
F statistic	$F_{(2,260)}$=0.98	$F_{(2,68)}$=0.91	$F_{(2,44)}$=2.40	$F_{(2,68)}$=2.90	$F_{(2,70)}$=1.31	$F_{(3,92)}$=320.25
p-value	0.3749	0.4058	0.1027	0.0616*	0.2764	0.0000***
Dependent variable – storage						
F statistic	$F_{(2,260)}$=2.26	$F_{(2,68)}$=1.99	$F_{(2,44)}$=0.02	$F_{(2,68)}$=1.18	$F_{(2,70)}$=0.98	$F_{(3,92)}$=340.59
p-value	0.1060	0.1453	0.9770	0.3135	0.3804	0.0000***

Notes: The p-values are for the F-test that testing that the coefficients of squares and cubes of the fitted values are jointly zero. *** denote rejection at the 1% level and * at the 10% level.

Thus, the market outcome, whether we look at price or storage, clearly differs from the optimal market outcome that would be possible within the experiment. In the following analysis the focus is on storage, partly because the results of the RESET test suggest that the linear model is appropriate overall and for all sessions but mainly because the experiment was framed as a storage experiment.

[2] A fixed effects model was chosen based on the results of a Hausman test.

Regression analysis is used to explore how carry-in stocks and harvest levels influence storage and whether or not history matters. The comparison of the results with those that would occur if all participants used the optimal strategy is complicated by the fact that the model is clearly non-linear when the optimal strategy is used but that it is linear in the experiment. Thus the linear model for the optimal strategy is known to be mis-specified.

For the results of the experiment, Table 6, therefore, reports the coefficients of a linear fixed effects regression where storage is regressed on carry-in stocks, harvest and the lagged price. These coefficients are compared to those based on the simulated optimal strategy for two models, firstly, the same linear fixed effects regression and, secondly, a Tobit model, which is the more appropriate specification for the optimal strategy.

The coefficients for carry-in stocks are similar for the two linear models. Thus, carry-in stocks had a similar average effect on storage in the experiment as would be the case if the optimal strategy had been followed. However, while in the experiment the linear model with the constant marginal effect is the appropriate specification, this is not the case for the optimal strategy. The marginal effect in the optimal strategy is not constant. The marginal effect is zero for the censored observations and higher than the marginal effect in the linear model for the non-censored observations.

Table 6 Estimated coefficients in the experiment and for the optimal strategy

	Constant	Harvest	Carry-in stocks	Lagged price
Experiment				
coefficient	-11.8581***	0.1983***	0.5543***	1.2241***
t-statistic	-11.90	15.59	12.51	3.54
Optimal strategy - linear model				
coefficient	-22.9549***	0.4371***	0.5137***	-0.5961
t-statistic	-13.00	24.90	9.78	-0.66
Optimal strategy – Tobit model				
coefficient	-39.1036***	0.6644***	0. 7574***	0.3661
t-statistic	-28.30	40.14	24.69	0.75

Notes: The dependent variable is storage. The t-test statistics are reported testing that the null hypothesis that the coefficients is zero against a two-sided alternative. *** denote rejection at the 1% level and ** at the 5% level.

By contrast, the coefficient on harvest for the linear models is less than half in the experiment than that of the optimal strategy. An increase in the harvest of one unit on average only increased storage by 0.19 units; this compares with an average 0.44 unit increase if the optimal strategy is followed and an average 0.66 unit increase for the optimal strategy in the interval where storage is not censored. The coefficient on the lagged price is statistically different from zero in the linear model based on the results of the experiment but is not so in the models based on the results of the simulated optimal strategy.

In a market populated by stockholders behaving according to the competitive storage model, the decisions about storage depends exclusively on availability in the current period and the relationship between storage and availability would be non-linear due to censoring. The average sensitivity of storage to carry-in stocks is close to optimal but this average masks the underlying differences of these averages. The average sensitivity of storage to the harvest outcome is less than would be optimal and history in form of the lagged price only has a statistically significant impact in the experiment but not for the optimal strategy. The results from the experiment clearly show that actual behaviour differs markedly from the behaviour suggested by the competitive storage model.

6 Discussion and conclusion

The main results of the experiment are that, at the market level, prices vary more than optimally because mean storage is lower than optimal and the storage level varies less than is optimal. The results are similar to those of savings experiments, which often find that participants in laboratory experiments under-save (see e.g. Brown et al., 2009 for of results from savings experiments). Savings and storage models are formally similar (Gouel, 2013). Unlike the storage experiment, though the savings experiments investigate individual decisions while the main interest of this storage experiment is at the market level. In this respect, the storage experiment is similar to multi-period asset market experiments where the main interest is also at the market level.

Participants in the storage experiment were mainly students as is the case in most experimental studies, including multi-period asset market experiments and the storage experiment on the Zambian maize market. This raises the question if such a convenience sample of student participants biases the results, especially because the set-up of the experiment, even in this simple form, is still relatively complex.

To make sure that participants had understood the main features of the experiment, the experiment was not started until all participants had answered the control questions correctly. Most participants answered the control questions correctly in a short period of time suggesting that the general features of the experiment were clear to the majority of participants. The rating of the difficulty of the experiment by the participants after the last round of the experiment was 6, on a scale from 1 (very easy) to 10 (very difficult). One participant rated the experiment 10 and five 9. Almost half rated it 7 or 8, suggesting that the experiment was somewhat difficult but not extremely so for most participants.

Without running the experiment using real world traders, the question of whether using a student population biases the results cannot be answered definitively. However, results for other experiments have shown that using student participants are similar results to those using a wider population (e.g. Andersen et al., 2010). In asset market experiments, bubbles were also found in experimental markets using participants who are experienced in real financial markets (Gerding, 2007).

The extent to which these results may generalise also depends, to some degree, on the impact of the design of the experiment on the storage level and variability. Since this is the first experiment of this kind and no comparison with other designs is possible, it is difficult to draw any firm conclusions. A number of possible impacts should be mentioned though. Participants can only buy, sell and store one unit. The optimal strategy, to which actual behaviour is compared, takes into account the fact that only changes in storage in integer steps are possible.

If some participants do not behave optimally, as is clearly the case in the experiment, the design adopted in this experiment limits the response by other participants to this deviation from optimal behaviour. For example when the optimal strategy would lead to storage at the maximum level and some participants do not store anything, within the experiment the storage level will be too low. Other participants who might get to understand that generally the storage level is too low when availability is very high have no possibility to compensate by storing more. This will be the case in any design with a maximum storage capacity at the market level that is the sum of individual storage capacity limitations. It would not be the case if there was a maximum storage capacity at the market level but no individual maximum capacity limits e.g. when unused storage space could be transferred from one participant to another. In that case a participant who notices that generally storage is too low at very high availability could use his own storage capacity and that of other participants. In this way, the maximum at the market level could be reached even if some participants do not behave optimally. In the real world, storage capacity for grains is rarely (or never) exhausted. This means that in the real world storage might get closer to the optimal level when availability is large.

The auction design in the experiment was very simple and this could also have influenced the results. Each participant was required to submit a price in each round which might actually have led to more transactions – and thus to more storage variability – than would be the case in a continuous auction design. In the real world, stockholders do not have to submit prices in each period but can just stay out of the market. In the experiment, a straightforward strategy for a seller who does not want to participate in the market was to submit a price that is impossible within the experiment. Prices that sellers could submit were not limited and any price above 2.40 ECU was impossible and thus meant that the seller would not make any transaction in that round – an outcome equivalent of not participating in a continuous double auction.

The competitive storage model, used as benchmark in this study, assumes that stockholders are risk-neutral. However, experiments have shown that in general participants in laboratory experiments are not risk neutral (e.g. Harrison and Rutström, 2008; Holt and Laury, 2002). In other models isoelastic utility functions of the form $U(c) = \frac{c^{1-\rho}}{1-\rho}$ are often used which have the characteristic of constant relative risk aversion. Holt and Laury (2002) give ranges for the coefficient ρ for risk loving, risk neutral and risk averse agents. They suggest slightly risk averse agents have ρ between 0.15 and 0.41, risk average agents between 0.41 and 0.68 and very risk averse agents between 0.68 and 0.97. The derivation of the solution to the competitive storage model using constant relative risk aversion is complicated. Including constant absolute risk aversion is slightly less complex but still complicates the results significantly. Empirical and experimental studies tend to be more supportive of the constant relative risk aversion assumption than the constant absolute risk aversion assumption (e.g. Chiappori and Paiella, 2001; Levy, 1994; Szpiro, 1986). Hence, a benchmark with constant absolute risk attitude might not provide a much better benchmark than a benchmark with risk-neutral agents.

Risk averse stockholders will store less than the risk-neutral ones. Risk-aversion therefore is a possible explanation why mean storage is below the optimal. However, there are three factors that suggest that the impact of risk-aversion is probably small. Firstly, the stakes in this experiment were relatively small and risk aversion is generally smaller with small stakes (Holt and Laury, 2002). Secondly, the benchmark optimal strategy is already mildly risk averse because buyers round-down and sellers round-up the optimal price. Thirdly, risk averse participants had a simple strategy within the experiment – never buy and walk away with the initial endowment of 10 ECU. This simple strategy is to input a price of 0 in each round. This is an obvious simple strategy and requires only a minimal understanding of the experiment. None of the 88 participants followed this strategy. Only 2 participants out of the 88 never stored and they did not obviously follow a "no buy" strategy. The prices they submitted were not sufficiently low to guarantee that they would not buy and they varied the prices throughout the experiment. However, an improvement on the benchmark, or at least the elicitation of the level of risk aversion in future experiments, would be desirable.

Similarly, a more detailed analysis of the impact of the relatively small possible gains from storing would shed light on the robustness of the results. The small earning potential is in line with the model which assumes that profits should be eliminated if storage is below its maximum level. The fact that only two participants never traded suggests that the small differences in earnings might not have had a substantial detrimental effect on the trading incentives. Possibly participants understood that, though they do not make large profits by storing if others also store, they could make more substantial if others do not store at high availabilities. The experimental design could be adjusted to investigate if the size of the profit opportunities that exist within the optimal strategy changes storing behaviour.

Further research is needed to assess the robustness of these results with regard to a number of other characteristics of the experiment. The robustness to different harvest distributions is one obvious line of research. Different distributions might mean more or less diversion from optimal strategies. Different crops vary in the size of the yield variance and with more or less variance in the harvest shocks results might differ. Also, although, the three point distribution has the advantage of being simple and easily communicable to participants, it is not very realistic and results might or might not be robust to more complicated harvest distributions.

Eventually, the most interesting extension will be to assess the impact of market interventions on storage levels. The fact that average storage is below its optimal level and storage does not vary as much as would be optimal does not necessarily mean that market interventions would improve the outcome. Experiments including market interventions aimed at bringing storage closer to the optimal level would be required to check on their effectiveness.

To conclude, the results of the experiment show that within the experimental setting mean storage and the standard deviation of storage are below their optimal level leading to price series that vary more than is optimal. Whether or not any policy intervention could bring storage closer to the optimal level was not investigated in this study. Also, the robustness of the results in more realistic and complex settings will have to be investigated in future experiments before drawing any general conclusions from this experiment.

References

Abbink, K., Jayne, T. S. and Moller, L. C. (2011). "The Relevance of a Rules-based Maize Marketing Policy: An Experimental Case Study of Zambia", *The Journal of Development Studies* **47** (2), 207-230.

Andersen, S., Harrison, G. W., Lau, M. I. and Rutström, E. E. (2010) "Preference heterogeneity in experiments: Comparing the field and laboratory"*Journal of Economic Behavior & Organization* **73** (2), 209-224.

Brown, A. L., Chua, Z. E. and Camerer, C. F. (2009). "Learning and Visceral Temptation in Dynamic Saving Experiments" *The Quarterly Journal of Economics* **124** (1), 197-231.

Chiappori, P.-A. and Paiella, M. (2011). "Relative risk aversion is constant: Evidence from panel data" *Journal of the European Economic Association* **9** (6), 1021-1052.

European Commission "Tackling the challenge of rising food prices - Directions for EU action", Communication from the Commission to the European Parliament, the Council, the European Economic and Social Committee and the Committee of the Regions COM(2008) 321 final, Brussels: European Commission, 20 May 2008.

Fischbacher, U. (2007). "z-Tree: Zurich toolbox for ready-made economic experiments" *Experimental Economics* **10** (2), 171-178.

Gerding, E. F. (2007) "Laws Against Bubbles: An Experimental-Asset-Market Approach to Analyzing Financial Regulation" *Wisconsin Law Review* **5**, 979-1039.

Gouel, C. (2013). "Comparing Numerical Methods for Solving the Competitive Storage Model" *Computational Economics* **41** (2), 267-295.

Gustafson, R. L. "Carryover Levels for Grains: A Method for Determining Amounts that Are Optimal under Specific Conditions", Technical Bulletin 1178, US Department of Agriculture, Washington: USDA, October 1958.

Harrison, G. W. and Rutström, E. E. (2008). "Risk aversion in the laboratory" *Research in Experimental Economics* **12**, 41-196.

Holt, C. A. and Laury, S. K. (2002). "Risk Aversion and Incentive Effects" *The American Economic Review* **92** (5), 1644-1655.

Levy, H. (1994). "Absolute and relative risk aversion: An experimental study" *Journal of Risk and Uncertainty* **8** (3), 289-307.

Noussair, C. N. and Tucker, S. (2013). "Experimental Research on Asset Pricing" *Journal of Economic Surveys* **27** (3), 554-569.

Organisation for Economic Co-operation and Development (OECD) "Rising Food Prices: Causes and Consequences", Working Document, Paris: OECD, 2008. Retrieved from http://www.oecd.org/dataoecd/54/42/40847088.pdf

Ramsey, J. B. (1969). “Tests for Specification Errors in Classical Linear Least-Squares Regression Analysis” *Journal of the Royal Statistical Society. Series B (Methodological)* **31** (2), 350-371.

Smith, V. L., Suchanek, G. L. and Williams, A. W. (1988) “Bubbles, Crashes, and Endogenous Expectations in Experimental Spot Asset Markets” *Econometrica* **56** (5), 1119-1151.

Szpiro, G. G. (1986) “Measuring Risk Aversion: An Alternative Approach” *The Review of Economics and Statistics* **68** (1), 156-159.

Wiggins, S. and Keats, S. (2010). “Grain stocks and price spikes”, Annex to HM Government “The 200708 Agricultural Price Spikes Causes and Policy Implications”, HM Government Global Food Markets Group, London, January 2010.