THE EFFECTS OF SEEMINGLY NONBINDING PRICE FLOORS: AN EXPERIMENTAL ANALYSIS*

STEPHEN SALANT†, WILLIAM SHOBE§, AND NESLIHAN ULER‡

Abstract. Price floors are common policies in markets for storable goods such as commodities, bankable emissions permits, and currencies. Hard price floors are implemented as unlimited government buybacks and prevent the price from falling below the floor; soft floors, whether implemented as limited buybacks or as reserve prices in emission permit auctions, allow the market price to fall below the floor. We specify and then test in the laboratory a two-period model with the same properties as our infinite-horizon model, Salant et al. (2022). Theory predicts that asset prices will respond to price floors even in a set of circumstances where the floor seems nonbinding. Most of our experimental findings are consistent with theoretical predictions: a seemingly nonbinding floor can cause the price and carryover to jump up, the jump is higher with a hard floor than a soft one, and when the floor is so low that the theory predicts no effect, none is observed. In contrast to theoretical predictions, however, the soft floor fails to increase the carryover and market price in our experiment.

JEL codes: C6, C92, D0, Q1, Q5

Keywords: Price floors, emissions markets, agricultural price supports, commodity markets, laboratory experiments

---

Date: May 27, 2023.

*An earlier version of this paper was circulated under the title, "The Effects of 'Nonbinding' Price Floors: An Experimental Analysis." The authors would like to thank Chris Barrett, Severin Borenstein, Dallas Burtraw, John Kagel, John List, Yusufcan Masatlioglu, Kyle Meng, and many seminar and conference participants for their valuable comments on previous drafts. Salant wishes to acknowledge useful discussions with Ping Han, Armon Rezai, Yichuan Wang at the inception of this research project. We gratefully acknowledge financial support from the International Foundation for Research in Experimental Economics (IFREE), the Erb Institute at University of Michigan and AREC at University of Maryland. We thank John Jensenius for his help with programming and running the experiment. We also thank Andrew Card, Ariel Listo, Aldo Gutierrez Mendieta and Haozhu Wang for superb research assistance.

† University of Michigan. Email: ssalant@umich.edu.

§ University of Virginia. Email: shobe@virginia.edu.

‡ University of Maryland and University of Michigan. Email: neslihan@umd.edu [Corresponding].

1. Introduction

According to common textbook presentations, a price floor set below the prevailing price should have no effects, but one set anywhere above that price should raise the price to the level of the floor. For example, whenever the minimum wage is below the prevailing wage, removal of the minimum wage should not affect the market wage.

Price floors on goods that can be stored —whether imposed on commodities, bankable emissions permits, or currencies — should not work like this. Private owners of these storable goods often carry part of the available stock into the future. The more they carry out of the current period, the higher will be the current price. How much they carry depends on a comparison of the current price to the expected discounted price in the next period. A price floor inserted strictly below the current price may nonetheless bind in some states in the future, requiring government purchases (or withholding of supplies). In that case, the floor will raise the expected price next period, stimulating carryovers, and widening the gap between the current price and the floor determining it. In Salant, Shobe and Uler (2022), we refer to the influence the price floor exerts on the price despite the gap between them as "action at a distance."[1] In fact, a gap between the first period price and the floor can arise even if the floor is inserted *above* the first period price since the additional carryover induced by the floor can raise the first-period price above the floor. Policy makers who observe a gap between the equilibrium price and the floor below may be surprised if they think the floor can be removed without causing the price to fall.

There are two types of price floors. If the floor is hard, the government commits to purchasing unlimited amounts at the floor price; if the floor is soft, the commitment is to purchase (or to refrain from selling) up to a specified maximum. Hard price floors have often been used to support the prices of agricultural commodities and currencies while soft floors are widely implemented as reserve prices, for example, in emission permit auctions.[2] The two floors affect the expected price differently. If the floor is hard, the price cannot fall below it, since the government can buy unlimited amounts; if the floor is soft, the market price can fall below the floor, since the government demand is limited.

[1] A similar phenomenon can occur with a price ceiling inserted above the current price. There the ceiling can lower the expected price, depress carryovers, and lower the current price so that a gap remains between the new current price and the ceiling above it.

[2] An auction reserve price is a special case of a soft floor in which the amount purchased at the floor price by the government is limited to the amount auctioned and the floor price equals the auction reserve price.

Therefore, as we show in Salant, Shobe and Uler (2022), action at a distance is expected to be stronger with a hard floor than with a soft floor.

Our companion paper (Salant, Shobe and Uler, 2022) investigates theoretically the effects of price floors on the price of storable goods in the standard infinite-horizon stochastic model of agricultural carryovers first analyzed by Gustafson (1958).[3] In this follow-up paper, we ask if action at a distance is a mere theoretical curiosity or if it describes actual behavior in response to price floors. To answer this question, we first build a simple two-period model (with the key properties that can give rise to action at a distance), which we then test in a laboratory experiment.

In the first part of the paper, we construct a simple two-period competitive model. Since the imposition of a floor (or ceiling) on the price of storable goods has been and continues to be such a popular policy, we could not hope to capture the diverse institutional settings in which it arises.[4] Instead we have retained only those elements that theory identifies as essential for action at a distance. We assume each agent has the same endowment of a storable good in the first period. Demand in the first period is known, as is the demand in each of the two possible states in the second period. Each agent knows the probability that each state will occur. In equilibrium, the aggregate carryover adjusts so that the price in the first period is equal to the discounted expected price in the second period. A price floor imposed below the first-period price may require the addition of government demand to private demand in the low-demand state of the second period. If so, it will stimulate additional carryovers from the first period and this will raise the first-period price. Moreover, the carryovers and the first-period price will be (weakly) larger under a hard floor than a soft floor.

[3]Subsequent contributions were made by Samuelson (1971), Kohn (1978), Gardner (1979), Newbery and Stiglitz (1981), Wright and Williams (1982), Salant (1983), Scheinkman and Schechtman (1983), Cottle and Wallace (1983), and Deaton and Laroque (1992), among others.

[4]The impulse to avoid the extremes of high or low prices of storables has resulted in the frequent use of buffer stock policies. Numerous proposals have been made for commodity price-stabilization with buffer-stock programs. Only political divisions prevented implementation of UNCTAD's "Integrated Program for Commodities" proposal to impose price bands on 17 individual commodities. The European Monetary System, which prevailed for two decades (1973-1992), attempted to keep foreign exchange rates within a "target zone"; the Chinese continue to keep its currency (the yuan) within one. The US Regional Greenhouse Gas Initiative and its counterpart in California (AB-32) have emission permit auctions with price floors and ceilings. In this paper, we investigate several properties of price floors that are equally applicable to the prices of commodities, currencies, bankable emissions permits, or any other storable asset created in the future.

Switching from the infinite-horizon formulation of our companion paper (Salant, Shobe and Uler, 2022) confers several advantages. Since the model terminates after the second period, nothing will be carried out of that period. This makes straightforward calculation of the state-dependent price in the second period as a function of the carryover. For any given carryover, the expected discounted price in the second period is simply the probability-weighted average of the state-dependent discounted prices. We construct a two-period parameterized example with two states in the second period and linear demand curves in both periods, relegating its nonlinear generalization to Online Appendix A. Both the example and its nonlinear generalization have the same qualitative properties as the infinite-horizon model in Salant, Shobe and Uler (2022). We use the example to deduce closed-form quantitative predictions, which we then test using a laboratory experiment.

To test our theoretical predictions, we design an experiment with five treatments: a baseline with no price floor, a pair with a hard price floor and a soft price floor set at the same level, and a pair with price floors so low that they should have no effect. Consistent with our theoretical model, we observe in the laboratory action at a distance when a seemingly nonbinding hard floor is imposed. We also find that the jump in the carryover (or in the initial price) is higher with a hard floor than with a soft floor. Finally, when a very low price floor is imposed, neither a soft nor a hard floor increases the carryover or initial price. However, contrary to our theoretical predictions, a soft floor does not affect the carryover or the market price. By conducting an additional follow-up experiment, we confirm that this experimental result is independent of whether the soft floor is implemented as a reserve price or as limited government buyback.

Previous studies have suggested the possibility of action at a distance in equilibrium models.[5] Using a deterministic Hotelling model, Dwight Lee (1978) showed that a hard price ceiling on exhaustible resources such as oil will affect today's price, even if the price would not hit the ceiling until well into the future. The "action-at-a-distance" hypothesis

[5]Related phenomena can also occur in single-agent models. Solow (1974), Dow and Olson (2001), and Olson (1989) provide examples where an optimizer alters his current behavior in response to a situation that is anticipated to affect him only in the future. Solow's central planner increases his current oil extraction when he discovers the availability of a high-cost perfect substitute even though he has no current use for this expensive "backstop" technology. Dow and Olson's consumer increases his current precautionary saving in response to a liquidity constraint which, although not currently binding, may bind with some probability in the future.

we test here is, in part, a stochastic generalization of Lee's insight about hard ceilings; Lee did not consider soft ceilings.[6]

By treating hard and soft floors in a unified way, Salant, Shobe and Uler (2022) shows that market price is at least as great with the hard floor as with the soft floor. Krugman (1991) examined the possible stabilizing effects of target zone policies in currency markets, but did not address action at a distance.[7] Schennach (2000) and Burtraw, Palmer and Kahn (2010, p. 4922) were the first to address the dynamic stochastic effects of price floors and ceilings in the emission market context. For a more detailed discussion of these papers, see Salant, Shobe and Uler (2022).

There is a growing experimental literature examining the effects of price floors or ceilings on market outcomes. Isaac and Plott (1981), and Smith and Williams (1981) used laboratory experiments to test the effects of price controls on the market price in the single-period (static) context. Other papers exploring the static context include Perkis, Cason and Tyner (2016), and Friesen et al. (2022). Multiperiod experiments involving price controls like our own began with Stranlund, Murphy and Spraggon (2014), who explore the effects of banking and price controls on price volatility. Cason, Stranlund and de Vries (2022) showed that a hard price floor increases investment in abatement cost reduction.[8] Neither paper studies the impact of a seemingly nonbinding price floor on the current price and carryover.

Holt and Shobe (2016) used experiments to explore some of the unintended consequences of the European Union Emission Trading System's (EU-ETS) Market Stability Reserve, a plan to raise emission allowance prices, not by directly changing supply, but rather by restricting the number of *banked* (carryover) allowances.[9] In their experiment on price collars (the combination of a soft ceiling and a soft floor), even when the reserve

---

[6]Hard ceilings, like hard floors, exert action at a distance. In a stochastic setting, soft ceilings, like soft floors, can be breached by the market price. But there the similarity ends. With a soft ceiling, every individual expecting any upward jump in prices when the ceiling is breached has an incentive to purchase an infinite amount, resulting in a speculative attack (Salant, 1983). With a soft floor, individuals expecting a downward jump in prices when the floor is breached are restrained from such extreme behavior by a non-negativity constraint, which limits sales to his holdings.

[7]For an exposition of the Krugman model and the follow-on research, see Svensson (1992) and the references therein.

[8]MacKenzie (2022) provides a useful review of the emissions auction literature emphasizing price control mechanisms.

[9]This indirect path to adjusting supply leads to a number of unintended consequences, as noted in Perino et al. (2022).

price was strictly below the current spot price (and ultimately never binding), raising the auction reserve price raised the observed spot price. Hence, Holt and Shobe deserve credit as the first to report action at a distance in the laboratory. Since that was not the focus of their experiments, they did not explore the phenomenon, elucidate its underlying mechanism theoretically, or distinguish hard and soft floors.[10] Distinguishing hard and soft floors is important since they can have quite different effects both in theory and in laboratory settings.[11]

We proceed as follows. Section 2 presents a tractable two-period model and illustrates the consequences of imposing each type of floor with a parameterized example involving linear demand curves. Section 3 describes our experimental test of the main predictions of this two-period example and reports our findings. Section 4 concludes the paper.

2. Theory

To test the hypotheses suggested by our theoretical paper (Salant, Shobe and Uler, 2022) with its infinite horizon and uncountable number of states, we strip away all unnecessary complications. We reduce the number of periods to two, since that is the smallest number necessary in theory to generate the phenomena of interest. Similarly, we eliminate uncertainty in the first period and reduce the number of states of nature in the second period to two. Although agent heterogeneity is an important feature of many markets, action at a distance should occur even when agents are identical. So we assume agents are identical. Since agents then should have no incentive to exchange any of their holdings with each other, we omit secondary markets and proceed as if there is a single representative agent who takes the three prices as given. We examine the effects of a floor that is imposed strictly below the first period price. The floor affects the equilibrium if and only if, in the second period, a floor at the same level mandates adding government demand to private demand in the low demand state.[12] Since the floor plays no role in the first period (i.e., removing it does not change the theoretical predictions), we eliminate it.

We assume that the identical agents initially hold in aggregate A_0 units of endowment. Agents use part of it in the first period and carry the rest (x units) into the second

[10]There are many differences between our experiments and those in Holt and Shobe (2016). See footnote 24 for more details.

[11]Fell and Morgenstern (2010), Fell et al. (2012), and others have treated soft floors as if they were hard. Care should be taken to specify the correct policy when analyzing floors since the effects of hard and soft floors differ.

[12]We use demand shocks in our model, but we could as easily have used supply shocks.

period. Agents know the demand schedule in the first period and in each potential state of the second period; but, when they make their first-period decisions, they do not know whether demand will be high (h) or low (l), which occur respectively with probability π_h and $1 - \pi_h$. We confine attention to the parameterized linear demand case on which our experiment is based. The nonlinear demand case is analyzed in Online Appendix A. By assumption, agents do not discount and are risk-neutral.[13] At the beginning of the second period, agents learn which state has occurred. They then participate in an auction where g units are offered for sale.

When the floor is hard, the government offers to buy back at the floor price as many units as are offered. This prevents the second-period price from falling below the floor. A soft floor can be implemented through the imposition of an auction reserve price, which is the typical method used in emission permit auctions or as a limited buyback, where the buyback amount is limited to the number of units auctioned. In what follows, we treat the hard floor case first and then explain how the soft floor case differs from it.

We define p^H as the price in the first period and p_i^H as the price in state $i = h, l$ in the second period when the floor is hard (denoted by superscript H). Assume the inverse demand curves are linear with the same slope but different intercepts: $p = m\{a_i - Q\}$ where $i = 0, h, l$. The demand parameters are $m = 2,\ a_0 = 100,\ a_h = 130,\ a_l = 70,\ \pi_h = .5$. The supply parameters are $A_0 = 40$ and $g = 8$. We assume the bidding in the auction is competitive.[14]

In equilibrium, the aggregate carryover (x) adjusts so that the price in the first period equals the price expected in the second period: $p^H = 0.5p_h^H + 0.5p_l^H$. Using the parameters of the model, we get the following:

$$(1) \qquad 2\{100 - (40 - x)\} = (.5)\ 2\{130 - (8 + x)\} + .5\max(f, 2\{70 - (8 + x)\}).$$

[13] The case of risk-aversion is discussed in Online Appendix B.

[14] When the floor (hard or soft) is implemented as a buyback, the bids and resulting auction price play no role in our theoretical predictions, since g units are always sold at auction. They also play no role when the soft floor is implemented as an auction reserve price provided agents in aggregate purchase g units at auction when the state-contingent price exceeds the auction reserve price and zero units when the state-contingent price is below the reserve price.

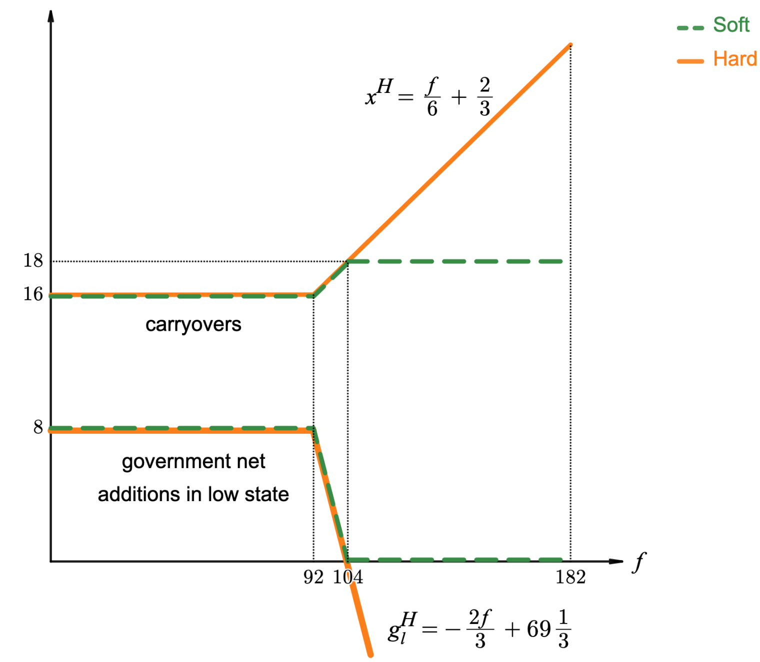

(A) Carryovers and Net Additions

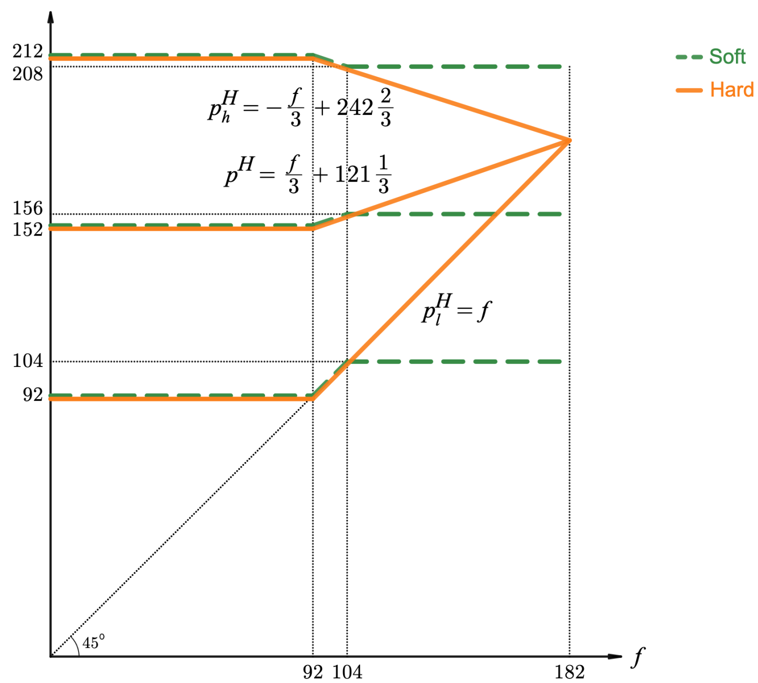

(B) Prices

FIGURE 1. Graphical Representation of the Model for the Parameters Used in Our Experiment

Solving for x, we obtain: $x = \max(16, \frac{f}{6} + \frac{2}{3})$. Thus, for $f > 92$, x strictly increases linearly in f. Substituting x into each term in braces, we obtain:

$$p^H = \max(152, \frac{f}{3} + 121\frac{1}{3}) \tag{2}$$

$$p_h^H = \min(212, 242\frac{2}{3} - \frac{f}{3}) \tag{3}$$

$$p_l^H = \max(92, f). \tag{4}$$

Note that if $f \leq 92$, $p_h^H = 212, p_l^H = 92$, and $p^H = 152$, the probability-weighted average. The reason p_h^H strictly decreases for $f \in (92, 182)$ is that the carryover strictly increases with f and, since no government intervention occurs in the high state, a demand-inducing price reduction is required to absorb the increased supply. If demand is low, that same increase in carryover would drive the price below f if the government did not intervene. Therefore, the government intervenes by adding fewer than 8 units to the market. With a hard floor, it does so by selling 8 units in the auction and then buying as much as is offered at the price f. If it purchases more than the 8 units auctioned, government net additions in the low state (denoted g_l^H) are negative—the amount bought by the government exceeds the amount auctioned.[15]

If $f \leq 92$, the government does not support the floor even when demand is low and so $g_l^H = 8$. If $f > 92$ it reduces g_l^H to raise the market price to the floor:

$$f = 2\{70 - (g_l^H + (\frac{f}{6} + \frac{2}{3}))\} \implies g_l^H = \min(8, 69\frac{1}{3} - \frac{2f}{3}), \tag{5}$$

which becomes negative if $f > 104$.

In the case of a soft floor implemented as an auction reserve price, the government intervenes by selling fewer than 8 units in the auction. Moreover, the government net additions must be nonnegative $g_l^H \geq 0$. When $f = 104$, net additions are zero and $x = 104/6 + 2/3 = 18$. Therefore, when $f > 104$, the effects of hard and soft price floors differ. When the soft floor is raised above 104 government net additions remain zero, the carryover remains 18, and the three prices do not change. Hence, unlike the hard floor case, the first period price can fall below the price floor. Note that, in theory, the soft

[15] In this example, the equilibrium price with no floor is 152. But as Figure 1 (b) reflects, any floor $f \in (92, 182)$ raises the first-period price of the storable so that it strictly exceeds the floor. If instead the floor were imposed on the price of a nonstorable (labor services), the floor (minimum wage) would have no effect if set below 152 and and, if set above it, would raise the wage only to the level of the floor, never above it.

floor can equally well be achieved with a limited buyback. We summarize these results in Figure 1.

3. A Laboratory Experiment to Test the Theory

3.1. **Experimental Design and Procedures.** The simple two-period model that we presented in the previous section generates sharp predictions that we can test in the laboratory: the level of carryover may jump when a "nonbinding" floor is imposed and, given the height of the floor, any jump will be weakly larger with a hard price floor than with a soft one. Our experimental design is based on the parameterized example we provided in Section 2. Recall that our model predicts that, for an extremely low price floor, neither soft nor hard floor will affect the carryover. To test this prediction, we use a floor of 68. In contrast, action at a distance is predicted to occur for floors higher than 92. Since in our experimental design, we wanted the two floor policies to have distinguishable effects, we set the floor at 128 so that:

(1) it does not bind in the high-demand state
(2) it is larger than 104.

Table 1 provides a summary of the experimental design and gives our model's predictions. In total, our experiment comprised five treatments: BASELINE, SOFT_HIGH, HARD_HIGH, SOFT_LOW and HARD_LOW. Subjects were randomly distributed to different treatments. Each treatment consisted of six sessions with ten subjects participating.[16] Each subject participated in only one session.

In addition to these main treatments, we conducted a follow-up study. The follow-up study took place after the main experiment was completed, using identical procedures and experimental laboratory. It consisted of only one additional treatment, SOFT_HIGH_LB, which was intended to be a robustness check. In SOFT_HIGH_LB treatment, we implemented the high soft floor by using limited buyback instead of an auction reserve price. Theoretical predictions for this follow-up study are identical to

[16] In choosing this design, we performed an analysis to guarantee that our statistical tests would have enough power to detect a difference of at least two units of carryovers. In order to have a power of 0.8 with $\alpha = 0.05$, we calculated that three (six) sessions are needed to detect a two (1.25) units of difference. Note that this analysis was based on two-sided tests and standard deviation calculations using the data from Part 2 of the first session of each treatment (which was approximately equal to one). If one instead uses one-sided tests, then naturally the power of our tests would increase (or, alternatively, one could detect smaller differences keeping the statistical power at 0.8).

those of the SOFT_HIGH treatment. In what follows we focus on the main experiment, deferring discussion of the follow-up study and its findings until Section 3.3.

	Price Floor	Carryover x	Period 1 Price p	Period 2 Price if Demand is High	Period 2 Price if Demand is Low
BASELINE	–	16	152	212	92
SOFT_HIGH	128	18	156	208	104
HARD_HIGH	128	22	164	200	128
SOFT_LOW	68	16	152	212	92
HARD_LOW	68	16	152	212	92

Table 1. Predictions in a Two-Period Example

With a floor of 68, there are no differences in carryovers or prices among the treatments: no floor (BASELINE), soft floor (SOFT_LOW) and hard floor (HARD_LOW). With a floor of 128, the effects of the two floors differ: the carryover increases by two units and the price jumps to 156 in the soft floor treatment (SOFT_HIGH) and the carryover increases by six units and the price jumps to 164 in the hard floor treatment (HARD_HIGH).[17]

Given that the main aim of our experiment is to demonstrate action at a distance as well as to show the differential effects of the soft and hard floor policies, we could have run only three treatments (BASELINE, SOFT_HIGH and HARD_HIGH). Instead, we chose to run two more. The SOFT_LOW and HARD_LOW treatments allow us to identify potential behavioral biases regarding the introduction of a price floor, and provides a stronger test of our model. In particular, if subjects change their behavior when there is a price floor due to reasons not captured by our model, then these additional treatments might help us identify such biases.

Next we explain the procedures of our laboratory experiment. Our experiment took place at the Symons Hall Experimental Laboratory (SHEL) at the Department of Agricultural and Resource Economics at the University of Maryland in September-December of 2018. To program the experiment, we used the z-Tree experimental software (Fischbacher, 2007). Using the ORSEE software ((Greiner, 2015)), we recruited a total of 300 subjects from the University of Maryland students on a first-come-first-serve basis from a large pool of potential participants representing different majors and different grade levels. Subjects participated in sessions that last approximately an hour and a half including

[17]As can be seen in Figure 1 (A), maximum carryover under the soft floor policy is 18, and hence, our design gives the best chance for the soft floor to be able to increase the equilibrium carryover and price.

payments. The experimenter read the instructions aloud at the beginning of each session. Instructions to our experiment is provided in Online Appendix D.

At the beginning of the experiment, subjects were told that the experiment has two parts, and that the instructions for Part 2 will only be given after Part 1 of the experiment is completed. Before each part of the experiment starts, subjects were given a quiz that tests and reinforces student understanding of the experimental setup.

Each part has five rounds, ten rounds in total. In each round, subjects participate in a 2-period market environment based on the example we constructed in Section 2. Each subject has the role of a trader in a market for a generic commodity. We referred to the generic commodity simply by "grain." They are asked to decide how much "grain" to buy, sell, or store.

Part 1 of each session does not have any price floor and is identical among all our treatments. Part 1 serves two main purposes. First, it allows subjects to learn about the procedures of the experiment. Second, and more importantly, it lets us control for the unobservable characteristics of the "traders" in our experiment. This way we can effectively control for variations in the skill levels and other uncontrolled characteristics of participants (e.g., (Smith and Williams, 1981)).

Part 2 of a session differs across treatments. Our BASELINE treatment repeats the same 2-period market environment as in Part 1 for another five rounds. Other treatments change only one aspect of the environment relative to the BASELINE treatment. Hard floor treatments have a buyback guarantee at a set price in the spot market. Soft floor treatments have a positive auction reserve price.[18]

In every round, at the beginning of the first period, each of the 10 subjects receives four units of grain (since $A_0 = 40$ in the example in Section 2) and also E\$250. Then they participate in a spot market. In the spot market, they have the opportunity to sell some or all of the grain they own to (pre-programmed) *computerized* consumers.[19] In the first period, consumer demand for grain is $Q = 100 - 0.5P$ as in the example we provided in Section 2. To help subjects understand the prices for a given level of quantity, we

[18]Our main study implements a soft floor as an auction reserve price as is done in actual practice in emission markets. As we explain in Section 3.3., our follow-up study shows our results are robust to implementing the soft floor using a limited buyback policy.

[19]Using computerized consumers allows us to eliminate mistakes of "buyers" which might confound our results and to focus on the behavior of the "traders."

provided subjects with a table demonstrating the prices instead of giving them a formula (see Table D1, Online Appendix D).

As sellers in the spot market, participants are asked to post offers for the units of grain they would like to sell. For each of their units of grain, they may post a different offer. Any grain they do not sell in the first period is automatically carried into the second period. At the end of the second period, any unsold grain has no value. The market has a uniform clearing price determined by the intersection of bid (pre-programmed valuation) and ask arrays.[20]

In the second period, demand can either be high ($Q = 130 - 0.5P$) or low ($Q = 70 - 0.5P$) with equal probabilities. Similar to the demand function in Period 1, in order to make it easy for subjects, instead of giving formulas, we provide them with a table of the valuations of the computerized buyers (see Table D1, in Online Appendix D). Before we conducted this study, we randomly generated six different demand sequences to be used in each session. Therefore, while the subjects faced a different sequence of random shocks among different sessions, each treatment has the same six sequences. This design makes the treatments comparable.

At the start of the second period, subjects learn whether demand for grain is high or low. They can try to augment their stored inventories by bidding in a grain auction where exactly eight units of grain are available in a sealed-bid, uniform price auction; this conforms with the example in Section 2 where $g = 8$. Participants are able to bid (anonymously and simultaneously) for up to two units, subject to their bids not exceeding their current cash balance. The bids are ranked from high to low. The grain being auctioned is sold to the highest bidders at a uniform price that is the value of the highest rejected bid, so all winning bidders pay the same price. If there are eight or fewer bids, then all bids are successful and the grain is free. Any unsold grain is no longer available to the market.

[20]In order to determine the market price as well as who makes a trade, we rank offers from lowest to highest, giving any tied offers distinct (but randomly determined) ranks, and we rank valuations of computerized consumers from highest to lowest. Second, we pair the lowest offer with the highest valuation, the second lowest offer with the second highest valuation, until we run out of offers. Third, we discard incompatible pairs—pairs where the buyer values the unit at less than the seller requires to sell it. Every undiscarded pair results in a transaction. All transactions take place at the same price: the lowest undiscarded valuation. We refer to this price as the (spot) market price.

After the grain auction, subjects again participate in the spot market as sellers. Other than the change in the demand schedule, the rules of the spot market are the same as in the first period.

In our soft floor treatments, subjects are told at the beginning of Part 2 that there is one change compared to Part 1; there is a reserve price in the auction. Subjects are not able to submit any bids that are lower than the reserve price. Similarly, in our hard floor treatments, at the beginning of Part 2 subjects are told that there is one change compared to Part 1; there is a special computerized buyer willing to buy unlimited amounts at a certain price (our hard floor price) in the second period.

Upon completion of each session, one round from Part 1 and one round from Part 2 were randomly chosen for determining subject payment. Subjects' earnings in experimental dollars (E$) in these two rounds were added up and converted to US Dollars ($) at the following rate: 100 Experimental Dollars (E$) = 1.00 US Dollar ($). Subjects' earnings (plus the $7 show-up reward) were paid to them in private at the end of the experiment. Most subjects' earnings were in the range $20-$30. The mean earnings was $24.8.

Predictions of our two-period model lead to two testable hypotheses:

(1) The level of carryover (or the first period equilibrium price) in Part 2 is the same for BASELINE, HARD_LOW and SOFT_LOW.
(2) The level of carryover (or the first period equilibrium price) in Part 2 is smallest for BASELINE, followed by SOFT_HIGH and then by HARD_HIGH.

In Section 3.2 we report on tests of these two hypotheses and, therefore, focus on behavior in the first period. Since the market price in each period and the carryover between periods are linearly related, it does not matter for any of our statistical tests whether we examine prices or carryovers;[21] examining both in the text would be redundant. Since changes in carryovers is what drives the changes in prices in the two periods, we chose to focus on carryovers in the text. For readers interested in seeing how carryovers translate into prices, we provide a discussion on these in Online Appendix C. In addition, see Online Appendix C for experimental results regarding behavior in the second period (including the number of units auctioned, auction and market prices).

Our experiments test the comparative static predictions of our model by using a between-subjects comparison and relying mainly on data from Part 2 of each treatment.

[21]Note that the power calculations are identical also. Due to the linear relationship between these two variables, while the standard deviation in prices doubles in the data, the gap between predicted prices across treatments exactly doubles too (relative to the gap between predicted carryovers).

While a within-subject comparison is also possible, we intentionally avoid that, since it is vulnerable to "order effects." Nevertheless, data from Part 1 is still valuable. First, we use data from Part 1 to check whether the random distribution of subjects into different treatments is satisfied by comparing carryovers/prices observed in Part 1. As we will see in Section 3.2, the differences observed in carryovers or in the market prices across treatments are not significantly different in Part 1. Therefore, we can attribute the differences in carryover/prices we observe in Part 2 to treatment differences and not to different characteristics of traders in different treatments. Second, as a complementary analysis, we also study Part 2 behavior by controlling for behavior in Part 1.

3.2. **Experimental Findings.** This section presents our experimental results from our main treatments. For completeness as well as for ease of comparison, our tables and figures also include results from the follow-up study. Nevertheless, to keep the section focused we defer the discussion of results from the follow-up study to Section 3.3. Unless otherwise mentioned, all reported tests throughout the paper (including the Online Appendix) are two-sided.

Figure 2 shows average carryover in the experimental rounds for each treatment. The first five rounds correspond to Part 1 and the second five rounds correspond to Part 2. We see a general trend that, over the rounds, subjects learn to sell more units and to carry over less in Part 1. And consistent with random distribution of subjects into different treatments, behavior looks similar in Part 1 among different treatments. Formal testing of these observations are provided below.

In Part 2, carryover in the BASELINE treatment is very close to the theoretically predicted level of 16 in each round. In fact, Table 2 confirms that the mean carryover is exactly at the predicted level. Figure 2 clearly shows action at a distance in the treatment HARD_HIGH, and, consistent with our prediction, average carryover is highest in the HARD_HIGH treatment relative to any other treatment (including the SOFT_HIGH treatment). We also make an interesting observation. Recall that the introduction of a very low price floor should not have any impact on the carryovers. In contrast, looking at the treatments HARD_LOW and SOFT_LOW, the impact seems to be negative, if any. While we will establish below that the effect is not statistically significant at the 5% level, it is important to acknowledge that these two treatments demonstrate that the addition of a price floor does not automatically lead to higher carryovers.

Table 2 presents mean carryover from Period 1 to Period 2 for each of our treatments in both parts. We use session averages over all rounds for a given part as independent

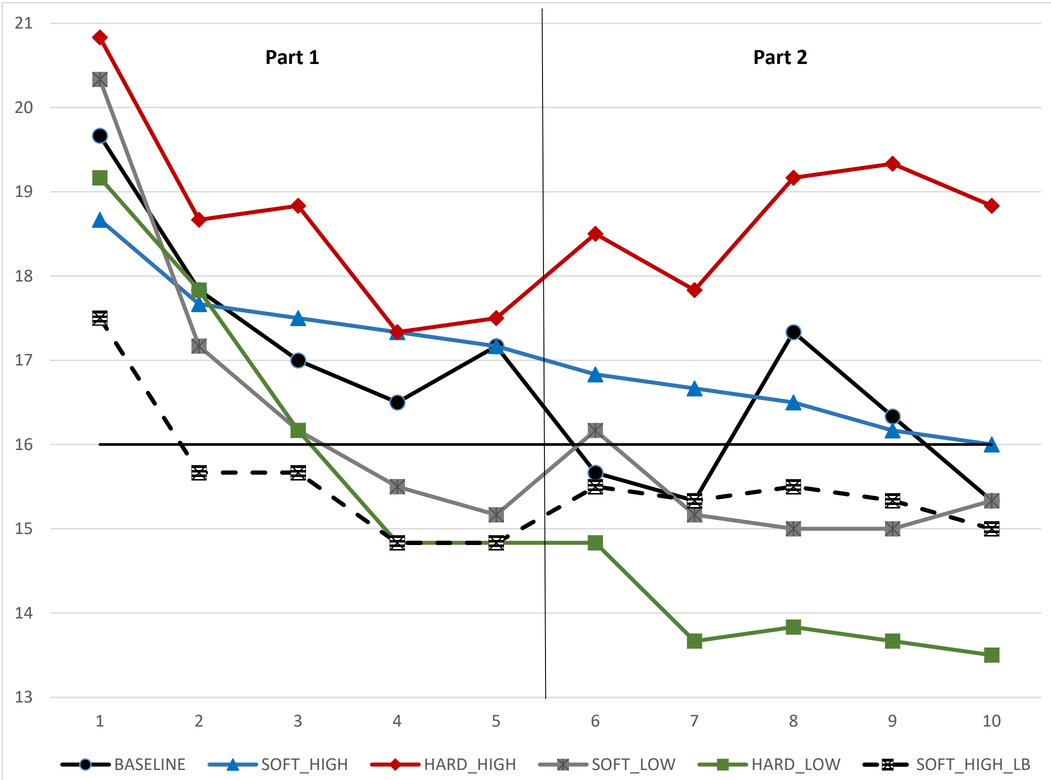

FIGURE 2. Average Carryover, All Rounds (5 in Part 1 and 5 in Part 2)

observations. To see whether differences in carryovers among different treatments are statistically significant in Part 1, we perform two-tailed, Mann-Whitney tests for each pairwise comparison (not shown here). As expected, none of the pairwise comparisons in Part 1 is statistically significant.[22]

The results of the Mann-Whitney tests for each pairwise comparison in Part 2 is presented in Table 3. Relative to the BASELINE treatment we see that the SOFT_LOW and HARD_LOW treatments are not effective, as predicted by our model—we observe no statistically significant differences among these treatments. These results are consistent with our Hypothesis 1.

Carryover in the HARD_HIGH treatment is significantly higher than carryover in all of the other treatments. While we report two-sided exact p-values in Table 3 to be conservative, we could rely on one-sided tests because our theoretical predictions regarding

[22]Most p-values are quite large with the smallest being equal to 0.12 (including the SOFT_HIGH_LB treatment).

	Part 1		Part 2	
	Observed	Predicted	Observed	Predicted
BASELINE	17.63 (0.72)	16	16.00 (0.60)	16
SOFT_HIGH	17.67 (0.70)	16	16.43 (0.64)	18
HARD_HIGH	18.63 (1.42)	16	18.73 (1.03)	22
SOFT_LOW	16.87 (1.17)	16	15.33 (0.79)	16
HARD_LOW	16.57 (0.75)	16	13.90 (0.96)	16
SOFT_HIGH_LB	15.70 (0.89)	16	15.33 (1.01)	18
Standard errors in parentheses.				

TABLE 2. Mean Carryover

	SOFT_HIGH	HARD_HIGH	SOFT_LOW	HARD_LOW	SOFT_HIGH_LB
BASELINE	0.554	0.056	0.558	0.212	0.926
SOFT_HIGH		0.089	0.411	0.082	0.558
HARD_HIGH			0.041	0.024	0.065
SOFT_LOW				0.258	0.937
HARD_LOW					0.333
Each cell reports exact p-values associated with two-tailed Mann-Whitney tests.					

TABLE 3. Are Carryovers Different Between Treatments in Part 2?

HARD_HIGH treatment are one-sided. In that case, naturally the significance level of these results would be higher. For example, when we use one-sided tests, the comparison between the HARD_HIGH and BASELINE (SOFT_HIGH) is statistically significant at 5% level with an exact p-value of 0.03 (0.04).

While the SOFT_HIGH treatment does show an increase in carryover compared to the BASELINE treatment, the difference is not statistically significant. In fact, Figure 2 shows that average carryover in SOFT_HIGH treatment converges to 16 units, which is the predicted level in the BASELINE treatment with no price floors.

Recall that our results on carryovers and prices are identical.[23] This implies Period 1 equilibrium price is statistically significantly higher in the HARD_HIGH treatment than in the BASELINE treatment, which is consistent with action at a distance (corresponding prices can be seen at Table C1 in Online Appendix C). Our results also imply the jump in price is larger for the hard floor than for the soft floor, which is again consistent with our theoretical predictions. Yet, in contrast to our prediction, the increase in the Period 1 price in the SOFT_HIGH treatment is not significant. These results, therefore, provide partial support for our Hypothesis 2.

We also conducted OLS regressions to check for robustness of our results (see Table 4). We regress individual carryover on treatment dummies and round. Specification 1 (2) considers carryover decisions in Part 1 (2), while Specification 3 repeats the same analysis as in Specification 2 but also controls for behavior in Part 1. As Figure 2 shows, at the end of Part 1, while all treatments converge close to the model's prediction of 16, there are nevertheless (statistically insignificant) differences across where treatments converge. Therefore, in Specification 3 we include an additional variable, Ind_carryover_part1_round5, which is equal to the amount of grain each subject carried over from Period 1 to 2 in the last round of Part 1. This allows us to effectively control for any differences in individual characteristics by taking into account different levels subjects converged to by the end of Part 1.

First, we see no statistically significant differences across treatments in Part 1. Second, in Part 2, in line with our hypotheses, individual carryover levels in the BASELINE treatment are not statistically different from those in the SOFT_LOW and HARD_LOW treatments. Third, our main results regarding the HARD_HIGH treatment are now even more statistically significant (relative to those reported in Table 3). As expected, action at a distance occurs in the HARD_HIGH treatment (p-value equals 0.018 in Specification 2 and equals 0.010 in Specification 3), i.e., individuals carry over significantly more to the second period in the HARD_HIGH treatment relative to the BASELINE treatment. Moreover, individual carryover is significantly higher in the HARD_HIGH treatment relative to all of the other treatments as well. In particular, subjects store more grain in the HARD_HIGH treatment relative to the SOFT_HIGH treatment (p-value equals 0.048 in Specification 2 and equals 0.026 in Specification 3). Finally, similar to the non-parametric analysis, OLS regressions do not show a significant effect of the SOFT_HIGH treatment on carryovers relative to the BASELINE.

[23]Therefore, Table 3 also provides the statistical results for prices.

Dependent Variable: Individual Carryover	Part 1 (1)	Part 2 (2)	Part 2 (3)
SOFT_HIGH	0.00	0.04	0.04
	(0.09)	(0.08)	(0.04)
HARD_HIGH	0.10	0.27*	0.26**
	(0.15)	(0.11)	(0.09)
SOFT_LOW	-0.08	-0.07	0.02
	(0.13)	(0.09)	(0.06)
HARD_LOW	-0.11	-0.21	-0.11
	(0.10)	(0.10)	(0.09)
SOFT_HIGH_LB	-0.19	-0.07	0.04
	(0.11)	(0.11)	(0.05)
Round	-0.08**	-0.01	-0.01
	(0.01)	(0.01)	(0.01)
Ind_carryover_part1_round5			0.45**
			(0.04)
Constant	2.00**	1.63**	0.86**
	(0.08)	(0.06)	(0.07)
N	1,800	1,800	1.800
R squared	0.01	0.01	0.22

Pairwise comparisons of SOFT_HIGH, HARD_HIGH, SOFT_LOW, HARD_LOW and SOFT_HIGH_LB in Part 1: None of the Wald tests are statistically significant at the 5% level (p-values range between 0.067 and 0.817).

Pairwise comparisons of SOFT_HIGH, HARD_HIGH, SOFT_LOW, HARD_LOW and SOFT_HIGH_LB in Part 2: (To simplify the exposition only the statistically significant results are reported below. The reported p-values correspond to the second specification. The ones reported in parentheses correspond to the third specification.)

H_0 : SOFT_HIGH = HARD_HIGH, p-value = 0.048 (0.026)
H_0 : SOFT_HIGH = HARD_LOW, p-value = 0.024 (0.091)
H_0 : HARD_HIGH = SOFT_LOW, p-value = 0.008 (0.025)
H_0 : HARD_HIGH = HARD_LOW, p-value = 0.001 (0.005)
H_0 : HARD_HIGH = SOFT_HIGH_LB, p-value = 0.015 (0.030)

Note: * indicates statistical significance at the 5% level and ** at 1%. Robust standard errors clustered at the session level in parentheses. 36 clusters in total.

TABLE 4. OLS Regression Analysis

3.3. **Discussion and a Follow-up Study.** Our main experiment shows strong evidence for action at a distance when the floor is hard and above a certain threshold, and its effect is stronger than with a corresponding soft floor. In addition, consistent with the model, we see no evidence of action at a distance when the floor is sufficiently low. However, contrary to the prediction of the model, we fail to see evidence of action at a distance in the SOFT_HIGH treatment. It is puzzling to find that a soft floor is ineffective in our experiment. In this subsection, we investigate possible reasons for this deviation from our hypothesis.[24]

We can eliminate the possibility that the lack of response to a soft floor is attributable to the failure of subjects to understand the instructions for this particular treatment. Recall that the SOFT_HIGH treatment differed from the baseline treatment in only one respect: Subjects were told, at the beginning of Part 2, that no bids lower than the auction reserve price may be submitted in the second year. We enforced this by the computer automatically refusing to register lower bids. Moreover, at the beginning of each part, for each session, subjects completed a quiz to check their understanding, received an automated explanation of the correct answer on their computer screens if they made a mistake and had a chance to ask questions to the experimenter, if any, before the experiment started.

We then investigate whether the lack of action at a distance observed in this treatment is related to mistakes subjects made bidding with a reserve price in the low-demand state of the second period, which in turn could be due to insufficient experience with reserve-price auctions when making no purchases is optimal.

In the SOFT_HIGH treatment, our theory predicts that there should be no bids in the auction if there is low demand in the second period since the market price will be lower than the reserve price of 128. As we show in Figure C.1 in Online Appendix C, subjects purchased approximately 5 units on average in these auctions when the demand was low. Subjects lost money on every unit they bought in such state-contingent reserve price auctions. In particular, they paid 128 per unit at the auction but the subsequent market price they experienced was only 106.78 per unit (averaged across the six sessions) as shown in Table C.3 (with a minimum of 100 and a maximum of 118). Presumably subjects experiencing such losses would learn to buy less and, eventually, to buy nothing. Subjects in each session do buy successively less on each occasion when a low demand is

[24] Although the soft floor in Holt and Shobe (2016) was effective while ours was not, there are too many differences between the two experiments for a comparison to be illuminating. They used a 30-period session with a pre-announced, declining allowance allocation over time. With unlimited banking, the declining path of allowances gave subjects substantial incentive to smooth allowance costs over time. The soft floor (auction reserve price) was constant across periods. Subjects could anticipate the possible future cancellation of unsold permits, which would randomly reduce the long-run supply. Holt and Shobe's setup also included a soft price ceiling which would, if triggered, have resulted in the release of up to a limited number of additional permits. In our experiment, subjects play a two-period market game repeatedly and gain experience over repetitions in a given session.

realized, but the learning of subjects is incomplete. Recall that no subject in these six sessions with the SOFT_HIGH treatment experienced more than three auctions where it was optimal to bid nothing; but each experienced more than twice as many auctions (five in Part 1 and at least two in Part 2) where it was optimal to acquire all 8 units in the auction.

Since net additions in the low-demand state were much larger than zero (the theoretical prediction), subjects *anticipating* these larger net additions in the low-demand state would rationally carry less over to the second period. Therefore, it is possible that the reason we see no effects of a soft floor on the carryover and market price is the additional layer of difficulty a reserve price induces when subjects are deciding on the optimal course of action.

This additional layer of difficulty disappears if a soft floor is implemented as a limited buyback instead of a reserve price. While implementing a soft floor with a reserve price is *theoretically* equivalent to a limited buyback, these mechanisms might not be *behaviorally* equivalent. To examine whether a soft floor implemented with a limited buyback is effective, we also conducted the follow-up study, SOFT_HIGH_LB treatment. The SOFT_HIGH_LB treatment, differs from the HARD_HIGH treatment in only one respect: the buyback is limited to eight units, precisely the number of units being auctioned. The auction in the SOFT_HIGH_LB treatment has no reserve price. Instructions for this treatment are included in Online Appendix D.

The follow-up study was conducted in March 2023 in the same experimental laboratory at the University of Maryland as the main experiment using identical procedures. Like the design of the main experiment, the follow-up study consists of six sessions with exactly 60 subjects.

The Mann-Whitney tests comparing the carryover in Part 1 in the SOFT_HIGH_LB treatment to those of the main treatments reveal no statistically significant differences (all p-values are larger than 0.12). This is reassuring, especially because this experiment was conducted after the main experiment was completed.

To see whether the way the soft floor is implemented matters, we now compare the SOFT_HIGH and SOFT_HIGH_LB treatments. Table 3 shows no statistically significant difference in carryovers in Part 2 either. Moreover, a regression analysis reported in Table 4 with and without controlling for past behavior also finds no difference between these two treatments (p-values are 0.33 and 0.91, respectively). Therefore, we conclude the two theoretically equivalent methods of imposing the soft floor generate the same outcomes in our experiment.

To provide further support for our main results, we also compare the carryovers in SOFT_HIGH_LB to the BASELINE and the HARD_HIGH treatments. As with our main results reported in Section 3.2, we find no difference in carryovers between the BASELINE and SOFT_HIGH_LB treatments, while the carryover in the HARD_HIGH treatment is significantly higher than the SOFT_HIGH_LB treatment (Tables 3 and 4).

Our follow-up study permits us to eliminate another possible explanation for why a soft floor is ineffective in our experiment. Naturally, given the results in Holt and Shobe (2016), we are not saying that soft floors are always ineffective. We are, however, saying that a soft floor might not always be effective even when the theory predicts it will be. Future research is required to investigate this issue further.

4. Conclusions

In this paper, we constructed a two-period model to show why the price of storable goods (commodities, bankable emissions permits, or currencies) may be determined by a price floor even though the price floor lies strictly below that price. Similarly, a price ceiling may exert a downward force on the price even though the ceiling lies strictly above it. Floors raise the expected future price, stimulating carryovers which drive up the first-period price; ceilings lower the expected future price, depressing carryovers which drive down the first-period price. We showed theoretically the circumstances when hard and soft floors set at the same level should generate the same gap between the price and floor and when a hard floor should generate a strictly larger gap.[25] The mechanism we identify likely explains the results observed by Holt and Shobe (2016) in their laboratory experiments on EU-ETS Market Stability Reserve.

Using linear demands, we derived closed-form results more suitable for testing than the infinite-horizon formulation of Salant, Shobe and Uler (2022).

Policies aimed at limiting prices of storable assets need to take into account the current effects of policy regimes that are seemingly nonbinding before they are imposed or removed. Policy makers and analysts need to recognize the distinction between hard and soft price floors since, as we have clarified, their effects are often different. Finally, policy

[25]Our analysis justifies the conjecture in the *Financial Times* (2013) regarding the hard floor in the Chinese rice market. In that newspaper article, it is asserted that the hard floor affects the market price even though it lies significantly below that price.

designers can consider replacing an auction reserve price at f with a limited buyback at f. This would give them more flexibility since policy makers would no longer be constrained to set the maximum buyback equal to the amount sold at auction.

Our laboratory experiments confirm most of the theoretical predictions. For both soft and hard floors, if the floor is low enough, then the floor does not change carryover and the equilibrium price, but as the floor rises, it begins to push up the carryover and equilibrium price in the first period. In addition, consistent with our model, we are able to reject the hypothesis that there is no difference in the effects of soft and hard floors. Contrary to our theoretical prediction, however, we fail to observe action at a distance with a soft floor regardless of whether it is implemented as an auction reserve price or as a limited government buyback. This result is inconsistent with the pattern observed in Holt and Shobe (2016). This unexpected result will be a fruitful area for future research.

References

Burtraw, Dallas, Karen Palmer, and Danny Kahn. 2010. “A Symmetric Safety Valve.” *Energy Policy*, 38: 4921–4932.

Cason, Timothy N., John K. Stranlund, and Frans P. de Vries. 2022. “Investment Incentives in Tradable Emissions Markets with Price Floors.” *Journal of the Association of Environmental and Resource Economists*. Forthcoming.

Cottle, Rex L., and Myles S. Wallace. 1983. “Economic Effects of Non-binding Price Constraints.” *Journal of Industrial Economics*, 31(4): 469–474.

Deaton, A., and G. Laroque. 1992. “On the Behaviour of Commodity Prices.” *Proceedings of the National Academy of Science*, 59: 1–23.

Dow, James, and Lars Olson. 2001. “An Analytic Solution to a Stochastic Consumption/Saving Problem with Liquidity Constraints.” *Southern Economic Journal*, 459–464.

Fell, Harrison, and Richard D. Morgenstern. 2010. “Alternative approaches to cost containment in a cap-and-trade system.” *Environmental and Resource Economics*, 47(2): 275–297.

Fell, Harrison, Dallas Burtraw, Richard D. Morgenstern, and Karen Palmer. 2012. “Soft and hard price collars in a cap-and-trade system: A comparative analysis.” *Journal of Environmental Economics and Management*, 64(2): 183–198.

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

Friesen, Lana, Lata Gangadharan, Peyman Khezr, and Ian A. MacKenzie. 2022. “Mind your Ps and Qs! Variable allowance supply in the US Regional Greenhouse Gas Initiative.” *Journal of Environmental Economics and Management*, 112: 102620.

Gardner, Bruce. 1979. *Optimal Stockpiling of Grain.* Lexington Books.

Greiner, B. 2015. “Subject pool recruitment procedures: organizing experiments with ORSEE.” *Journal of the Economic Science Association*, 1: 114–125.

Gustafson, Robert L. 1958. “Carry-Over Levels for Grains.” Washington: Department of Agriculture Tech. Bulletin 1178.

Holt, Charles A., and William M. Shobe. 2016. “Reprint of: Price and quantity collars for stabilizing emission allowance prices: Laboratory experiments on the EU ETS market stability reserve.” *Journal of Environmental Economics and Management*, 76: 32–50.

Isaac, R. Mark, and Charles R. Plott. 1981. “Price Controls and the Behavior of Auction Markets: An Experimental Examination.” *American Economic Review*, 71: 448–459.

Kohn, M. 1978. “Competitive Speculation.” *Econometrica*, 46: 1061–1076.

Krugman, Paul. 1991. “Target Zones and Exchange Rate Dynamics.” *Quarterly Journal of Economics*, 106: 669–682.

Lee, Dwight R. 1978. “Price Controls, Binding Constraints, and Intertemporal Economic Decision Making.” *Journal of Political Economy*, 86(2): 293–301.

MacKenzie, Ian A. 2022. “The Evolution of Pollution Auctions.” *Review of Environmental Economics and Policy*, 16(1): 1–24.

Newbery, David M.G., and Joseph E. Stiglitz. 1981. *The Theory of Commodity Price Stabilization.* Oxford University Press.

Olson, Lars J. 1989. “Stochastic Growth with Irreversible Investment.” *Journal of Economic Theory*, 47: 101–129.

Perino, Grischa, Maximilian Willner, Simon Quemin, and Michael Pahle. 2022. “The European Union Emissions Trading System Market Stability Reserve: Does It Stabilize or Destabilize the Market?” *Review of Environmental Economics and Policy*, 16(2): 338–345.

Perkis, David F., Timothy N. Cason, and Wallace E. Tyner. 2016. “An Experimental Investigation of Hard and Soft Price Ceilings in Emissions Permit Markets.” *Environmental and Resource Economics*, 63(4): 703–718.

Salant, Stephen W. 1983. “The Vulnerability of Price Stabilization Schemes to Speculative Attack.” *Journal of Political Economy*, 91(1): 1 – 38.

Salant, Stephen, William Shobe, and Neslihan Uler. 2022. “The effects of “nonbinding” price floors.” *European Economic Review*, 145: 104122.

Samuelson, Paul. 1971. “Stochastic Speculative Price.” *Proceedings of the National Academies of Sceince*, 68: 335–37.

Scheinkman, J., and J. Schechtman. 1983. “A simple competitive model with production and storage.” *Review of Economic Studies*, 50: 427–441.

Schennach, Susanne M. 2000. “The economics of pollution permit banking in the context of Title IV of the 1990 Clean Air Act Amendments.” *Journal of Environmental Economics and Management*, 40: 189–210.

Smith, Vernon L., and Arlington W. Williams. 1981. “On Nonbinding Price Controls in a Competitive Market.” *American Economic Review*, 71(3): 467–474.

Solow, Robert. 1974. “The Economics of Resources or the Resources of Economics.”

American Economic Review, 1–14.

Stranlund, John K., James J. Murphy, and John M. Spraggon. 2014. "Price Controls and Banking in Emissions Trading: An Experimental Evaluation." *Journal of Environmental Economics and Management*, 68: 71–86.

Svensson, Lars E.O. 1992. "An Interpretation of Recent Research on Exchange Rate Target Zones." *Journal of Economic Perspectives*, 6: 119–144.

Terazona, Emiko. 2013. "Beijing's Rice Subsidies Buoy Imports." *Financial Times*.

Wright, Brian D., and J.C. Williams. 1982. "The Economic Role of Commodity Storage." *Economic Journal*, 92(367): 596–614.

FOR ONLINE PUBLICATION

Online Appendix A: General Model

We assume that the identical agents initially hold in aggregate A_0 units of the initial endowment of the (storable) good. We assume that the quantity demanded is a downward-sloping function of the price. When first-period decisions are made, the demand in the second period is uncertain (due perhaps to unforeseen fluctuations in the weather or the level of economic activity).

Using a unified framework, we study the effects of imposing the policy of no floor (N), a soft floor (S) or a hard floor (H). Each policy can be regarded as an auction of g units with no reserve price followed by a state-dependent constrained buyback at exogenous price f. The buyback, denoted $R_k^i \geq 0$, for $i = N, S, H$ and $k = h, l$, is constrained to zero in the case of no floor, is unconstrained[26] in the case of a hard floor, and has an upper bound of the auctioned amount (g) in the case of a soft floor set to mimic an auction reserve price. To see why an auction of up to g units with a reserve price f is equivalent to an auction with no reserve price combined with a buyback of up to g units at floor price f, note that under both policies the government places g units on the market if the price strictly exceeds f and 0 units if the price is strictly below f. Hence, the two policies are equivalent and result in the same market-clearing price. For a summary of the timeline for our unified framework, see Figure A.1.

Net additions by the government in state $k = h, l$ under policy $i = S, H$ are denoted by $g_k^i = g - R_k^i$. In other words, government's net addition is equal to the amount it auctions less the amount it buys back. When the second-period price is higher than the floor, the net additions equal to g since government buys nothing back or equivalently, in the case of a reserve-price auction, the government sells all g units. The market price, denoted by p^i for $i = N, S, H$, never falls below the floor with a hard floor policy but can fall beneath the floor with a soft floor.

In the latter case ($p^S < f$), net additions are always zero ($g_l^S = 0$) regardless of the height of the floor. If the soft floor is implemented using a buyback, the government

[26]A soft floor with a sufficiently large maximum purchase would duplicate the effects of a hard floor. For example, a government offer to buy at floor price f up to $g + A_0$ (the maximum that could possibly be offered) has the same effect as a hard floor.

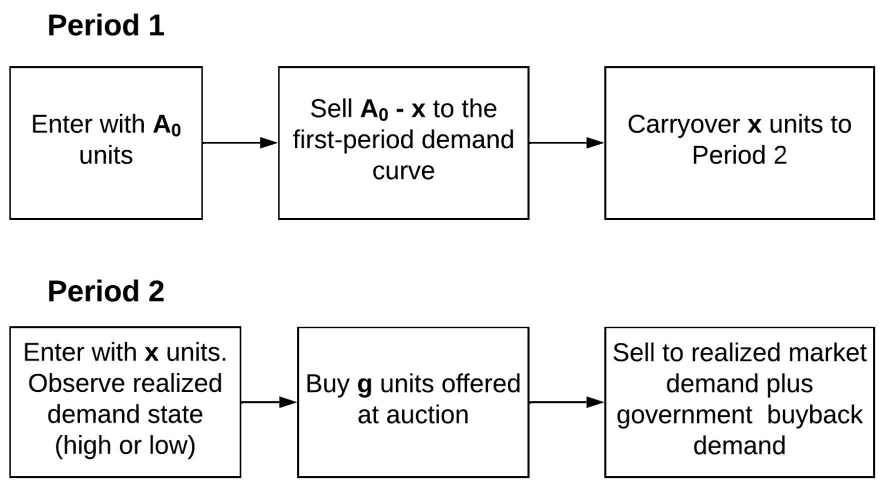

Figure A.1. In the second period, if the soft floor is implemented instead as an auction reserve price, the text in the middle box becomes “Bid for up to g units in reserve price auction” and the text in the right-hand box becomes “Sell to realized market demand.”

repurchases everything it auctions ($R_l^S = g$). If it is implemented as a reserve price, nothing is sold at the auction since bidder valuations are all lower than the reserve price.[27]

Under a soft price floor policy implemented as an auction reserve price or an equivalent buyback, the net addition can never be negative, while under a hard price floor policy the net addition can be negative.

Denote the inverse demand function in the initial period as $p(Q)$ and the inverse demand functions in the high and low states of the second period as $p_h(Q)$ and $p_l(Q)$ respectively, where Q denotes the quantity demanded. We assume these three functions are strictly decreasing and differentiable. $p(Q)$ represents the value of the marginal unit to the representative agent if he is already using Q units.

[27]The buyback at floor price f limited to g units and the auction of g units with reserve price set equal f generate the same net additions and the same prices. They do differ in one way which, although irrelevant in determining prices, is noteworthy nonetheless. In the buyback case, when the highest bidder valuation is below the floor price f, the government auctions at a lower price than it buys back. The counterpart of this capital loss of the government is the capital gain of those lucky enough to sell their part of g to the government buyback authority. No such capital gains and losses occur with the soft floor if it is implemented using an auction reserve price.

Since the shock resulting in high-demand would result in the same Q units having a higher value at the margin than would be the case in the first period (which in turn would have a higher marginal value than in the low-demand state), the functions can be ranked uniformly for any $Q: \; p_h(Q) > p(Q) > p_l(Q)$. We assume that, if nothing were carried over from the first period, the expected price in the second period would strictly exceed the first-period price. This assumption ensures that, in equilibrium, the carryover is strictly positive.

In equilibrium, agents will carry over just enough that the price in the first period will equal the price then expected in the second period. Formally, define the carryover under the *no floor* policy as the unique implicit solution x to $p(A_0 - x) = \pi_h p_h(x+g) + \pi_l p_l(x+g)$. Denote this solution as x^N.

With no floor, the equilibrium price in the high state strictly exceeds the equilibrium price in the low demand state since the price function in the high state lies uniformly above the price function in the low state and both price functions are evaluated in the no floor equilibrium at the same point $(x^N + g)$. That is, $p_h^N > p_l^N$, where p_j^N denotes the second-period price in state j $(j = h, l)$. Since the first-period price in the no-floor equilibrium (denoted p^N) is a probability-weighted average of the two state-dependent prices and $\pi_h \in (0, 1)$, it lies strictly between the two: $p_h^N > p^N > p_l^N$.

Imposing either a hard or a soft floor weakly *below* the smallest of these three prices (p_l^N) will have no effects. To see this, note that if a price floor below p_l^N were imposed, it would *a fortiori* be below both the first-period price and the price in the high-demand state. Since these three prices would all exceed the floor, no agent would sell to the government when he could earn more elsewhere, and the three markets would continue to clear as if no floor had been imposed.

Imposing either a hard or soft floor between the lowest of the three prices and the price under the no-floor policy results in action at a distance since it creates a disequilibrium in the low-demand state which can only be eliminated in equilibrium by a larger carryover. The larger carryover drives up the price in the first period. The disequilibrium arises because the government would be offering to pay strictly more in the low-demand state than what sellers would receive in the absence of a floor $(f > p_l^N)$.[28]

[28]In the case of soft floor being implemented as a reserve price, the disequilibrium arises because nobody has an incentive to bid in the auction, and the auction shuts down, i.e., net additions by the government become zero.

With $f > p_l^N$, either type of floor is high enough that the government is called upon to intervene in the low state by limiting its net additions ($g_l^i < g$ for $i = H, S$) because the buyback is strictly positive. As long as f is not so high that the floor binds even in the high state, net additions in the high-demand state will remain equal to g.

Propositions 1-4 below clarify when the hard and soft floors have identical effects, when their effects differ and in each case what these effects are. We relegate proofs to the Appendix.

In Proposition 1, we consider a situation where a hard floor binds in the low-demand state of the second period and ask how a marginal increase in that floor would affect the equilibrium. We prove that the carryover would increase, raising the first-period price and the expected price next period but lowering the price in the high demand state.

Proposition 1. Properties of a Hard Floor: *If $f \geq p_l^N$, increases in a hard floor increase the carryover ($\frac{\mathrm{d}x^H}{\mathrm{d}f} > 0$) and decrease the net additions in the low-demand state ($\frac{\mathrm{d}g_l^H}{\mathrm{d}f}$) < 0, with the magnitude of the decreases larger than the magnitude of the increases ($-\frac{\mathrm{d}g_l^H}{\mathrm{d}f} > \frac{\mathrm{d}x^H}{\mathrm{d}f} > 0$). As a consequence, the initial price must strictly increase, the price in the low state must strictly increase and the price in the high state must strictly decrease.*

Proof for Proposition 1. If $f > p_l^N$, the equilibrium carryover under a hard floor (x^H) and the net additions in the low-demand state (g_l^H) satisfies the following pair of equations:

$$\begin{aligned} &(1) & p(A_0 - x^H) &= \pi_h p_h(x^H + g) + \pi_l f \\ &(2) & p_l(x^H + g_l^H) &= f. \end{aligned}$$

Differentiating (1) with respect to f yields $\frac{\mathrm{d}x^H}{\mathrm{d}f} > 0$ and differentiating (2) yields $\frac{\mathrm{d}g_l^H}{\mathrm{d}f} < -\frac{\mathrm{d}x^H}{\mathrm{d}f} < 0$, where these derivatives are signed because each of the three inverse demand curves is strictly decreasing. We denote the solution to (1) and (2) as $x^H(f)$ and $g_l^H(f)$. □

Intuitively, if the carryover did not change when f increased, the expected price would strictly exceed the initial price, stimulating increased carryovers. An increase in the carryover raises the initial price and lowers the price expected to prevail in the second period, restoring the equality. In order for the price in the low-demand state to rise enough to match the increase in the floor, fewer units must be put on the market in the low-demand state despite the larger carryover; this occurs because the decline in the government's net additions more than offsets the increased carryovers, so their *sum* decreases.

As the hard floor is increased further, net additions in the low-demand state continue to fall until, for a unique floor (denoted f^*), net additions in the low-demand state reach zero. At that point, the amount the government purchases when supporting the hard floor equals the g units sold at auction: $g - R_l^H = 0$.

Proposition 2 shows that hard and soft floors implemented as an auction reserve price have identical effects for floor levels below f^*, the unique solution to $g_l^H(f) = 0$. Since g_l^H is a strictly decreasing function of f, $g_l^H > 0$ for $f < f^*$.

Proposition 2. When Soft Floors Induce the Same Responses as Hard Floors: *For $f \leq f^*$, soft and hard floors generate the same prices, carryovers, and net additions. Thus, if $f \in (p_l^N, f^*]$, hard and soft floors are equally active ($g_l^S = g_l^H \in [0, g)$) and if $f \leq p_l^N$, they are equally inactive ($g_l^S = g_l^H = g$).*

Proof for Proposition 2. For any $f \leq f^*$ (that is, for $f \leq p_l^N < f^*$ or $f \in (p_l^N, f^*]$), set the carryover and the net additions in the low-demand state under the soft floor to be the same as the equilibrium values of these two variables under the hard floor policy: $x^S(f) = x^H(f)$ and $g_{l(f)}^S = g_l^H(f) \in [0, g]$, where $g_l^S(f)$ denotes net additions in the low-demand state in response to a soft floor f. Since net additions under the hard floor policy are non-negative, they are feasible under the soft floor policy. Since the three markets would continue to clear, the soft and hard floors generate the same equilibrium. □

Proposition 3 demonstrates that, for any floor strictly greater than f^*, carryover is strictly greater with the hard floor than with the soft floor and so the current price must also be strictly greater with the hard floor than with the soft floor. If a hard floor is increased above f^*, all of the changes described in Proposition 1 continue to occur. The carryover rises and with it the first-period price, the expected second period price, and the price in the low-demand state. When the hard floor is increased above f^*, government net additions continue to decline, becoming negative for the first time: that is, the purchases by the government to support the price in the low-demand state exceed what is sold at auction ($g_l^H = g - R_l^H < 0$). But under a soft floor implemented as an auction reserve price, the government is not allowed to purchase more than it auctions so net additions can never be negative. In this case, for $f > f^*$ the government is powerless to purchase more under a soft floor than it did when $f = f^*$.[29] For the soft floor, $f > f^*$, the price in the low-demand state remains at the level induced by f^* even though the soft floor is higher. The first-period price as well as the carryover and prices in the two demand states remain at the level induced by the floor of f^* and net additions remain zero.

Proposition 3. When Soft Floors Induce Different Responses than Hard Floors: *For $f > f^*$, increasing the price floor does not increase the carryover or prices under a soft floor and hence the effects of the two types of floors differ. For floors in this range, (1) $x^H(f) > x^S(f) = x^S(f^*)$ and (2) $p^H(f) > p^S(f) = p^S(f^*)$.*

Proof for Proposition 3. For $f > f^*$, net additions under the hard floor in the low state (g_l^H) are strictly negative since $g_l^H(f)$ is strictly decreasing (Proposition 1) and $g_l^H = 0$ when $f = f^*$. However, negative net additions cannot be achieved with a soft floor; hence, $g_l^S(f) = 0$ when $f > f^*$ unchanged from $f = f^*$. Since the government behavior is unchanged, $x^S(f) = x^S(f^*) = x^H(f^*)$ for any $f > f^*$. The carryover and hence the three prices are independent of f since government purchases are the same for all floors higher than f^*. Since, as shown in Proposition 1, the carryover with a hard floor continues to increase with f in this range, the initial price if a hard floor ($f > f^*$) is imposed rises more than if a soft floor is imposed. □

If the soft floor is implemented as a buyback rather than as an auction reserve price, then the government is no longer constrained to limit its purchases to the amount auctioned. Our analysis is easily adapted for this policy as well.[30]

We now state our main proposition. First, either type of price floor inserted in a neighborhood on *either* side of the no floor equilibrium price (p^N) will raise the first-period price to a level above that price floor. This follows by noting (a) that either type of floor inserted at p^N will raise the first-period price above the floor since the floor must bind in the low state ($f > p_l^N$) and (b) a floor of either type inserted in a *neighborhood* on either side of p^N must also cause the first-period price to rise above the floor since $p^H(f)$

[29]If the soft floor were implemented instead as an auction reserve price, nothing would be sold at the auction.

[30]Suppose the lower limit on net additions is $g_{min} < 0$ instead of zero. Instead of f^*, there would be some $\hat{f}$ such that $g_k^H(\hat{f}) = g_{\min}$. The equilibrium under such a soft floor policy would coincide with the hard floor equilibrium for floors weakly smaller than $\hat{f}$ (the counterpart of our Proposition 2) but would differ for larger floors (the counterpart of our Proposition 3). In particular, for larger floors the net addition in the low demand state would remain $g_{\min}$ and the equilibrium under this soft floor policy would remain unchanged floors larger than $\hat{f}$.

and $p^S(f)$ are each continuous functions, being composite functions of two continuous functions: $p^i(f) = p\left(A_0 - x^i(f)\right)$ for $i = S, H$.

Second, if the first-period price under hard floor f^* exceeds that floor ($p^H(f^*) > f^*$), the same must be true of the first-period price under a soft floor. This follows since at f^* the two prices coincide. Moreover, for any f in a right-neighborhood of f^* the first-period price under a hard floor will strictly exceed the first-period price under a soft floor which is constant and must therefore continue to exceed the floor.

Proposition 4. Action at a distance: *If either type of floor is inserted in the neighborhood above or below the equilibrium initial price with no floor, the equilibrium initial price with the floor will strictly exceed the floor. Hence, the floor affects the initial price in equilibrium even though the gap between that price and the floor makes the floor appear "nonbinding." If in addition $p^H(f^*) > f^*$, then for floors in a neighborhood above f^*, $p^H(f) > p^S(f) > f$. That is, both types of floors will exhibit action at a distance with the gap between the initial price and the floor being strictly larger under the hard floor than under the soft floor.*

Proof for Proposition 4. Set $f = p^N$, the initial price under no floor. Denote this floor as f^N. Then, from Proposition 1, the floor will bind in the low state since $p_l^N < p^N = f^N$. Hence, $x^i(f^N) > x^N(f^N)$ for $i = H, S$ and the initial price under either type of floor must strictly exceed the *no floor* price and hence must strictly exceed the floor: $p^i(f) > f^N = p^N$ for $i = H, S$. Moreover, since the initial price under either type of floor is a continuous function of f (since it is a composite function of two continuous functions $p^i = p(A_0 - x^i(f))$ for $i = H, S$) any floor inserted in the *neighborhood* above or below f^N must also cause the initial price to exceed the floor.

If in addition $p^H(f^*) > f^*$, (for example, in Subsection 3 .2 $p^H(f^*) = 156$ and $f^* = 104$) then $p^S(f^*) = p^H(f^*) > f^*$. And since $p^H(f) > p^S(f)$ for $f > f^*$ and $p^S(f)$ is constant for $f > f^*$ there will be a right-hand neighborhood of f^* such that $p^H(f) > p^S(f) > f$.

□

Proposition 4 shows that there is a neighborhood around the first-period price in the no-floor policy such that action at a distance occurs.

Online Appendix B: Relaxing the Assumption of Risk Neutrality

One might ask whether our results are robust to dropping our risk-neutrality assumption. Under risk aversion the carryover in the absence of a floor will be smaller than under risk neutrality and so the initial price will be lower. But this initial price must still lie between the price in the two states; in particular, it must lie strictly above price in the low state; otherwise stockpilers would experience a capital gain whichever state is realized (a favorable one-sided bet). But then a floor of either type set equal to the initial price under no-floor policy will require intervention in the low state and the carryover must increase. Therefore, there will be a gap between the induced initial price and the imposed floor. Moreover, continuity of the carryover function associated with each type of floor under risk aversion implies that action at a distance would occur if the floor is inserted in a neighborhood on either side of the initial price under no-floor policy. Hence, all of our qualitative results continue to hold under risk aversion.

Under risk aversion, the three market-clearing equations (for the initial period and the two states of the second period) are unchanged. Moreover as before, when the hard floor binds in the low state, we have $p_l = f$. These four equations define four variables—the three prices and the net addition in the low state (g_l)—as functions of x and f just as in the case of risk neutrality. If, as is intuitive, an increase in the hard floor can be shown under risk aversion to lead to a strictly larger carryover, then market-clearing in the low state requires that the sum ($g_l + x$) of the net addition and the carryover falls even though the carryover induced by the increased hard floor rises. The same arguments used in the text to prove our four propositions under risk neutrality can then be repeated to extend these propositions to risk aversion.

To establish that under risk aversion the carryover is indeed strictly increasing in the hard floor when it binds in the low state, we assume that n stockpilers have identical preferences and each chooses his individual carryover y to maximize expected utility taking the prices as given. Hence, $x/n = \text{argmax}_{y \geq 0} \pi_h u(z + (p_h - p)y) + \pi_l u(z + (f - p)y)$, where z is the income each person would receive even if he carried over nothing and u is a strictly increasing, strictly concave von Neumann scaling function. It is then straightforward to verify that x is a strictly increasing function of f. We conclude, therefore, that assuming risk aversion would introduce complexity without changing any qualitative results. To keep the model as simple as possible, we assumed risk neutrality.

Online Appendix C: Additional Analysis

Table C.1 shows average prices observed in year 1 spot market both for Part 1 and Part 2. Session averages over all rounds for a given part are used, giving us six independent observations per treatment.

	Part 1		Part 2	
	Observed	Predicted	Observed	Predicted
BASELINE	155.27 (1.45)	152	152.00 (1.20)	152
SOFT_HIGH	155.33 (1.39)	152	152.87 (1.29)	156
HARD_HIGH	157.27 (2.84)	152	157.47 (2.05)	164
SOFT_LOW	153.73 (2.35)	152	150.67 (1.58)	152
HARD_LOW	153.13 (1.49)	152	147.80 (1.91)	152
SOFT_HIGH_LB	151.40 (1.77)	152	150.67 (2.02)	156
Standard errors in parentheses. There are 6 observations per cell.				

Table C.1. Mean Year 1 Prices

While our hypotheses are related to period 1 behavior, we also study Period 2 behavior, since it helps us understand subjects' choices in period 1. We start with the number of units purchased in the auction in Period 2. We expect subjects to buy all eight units in all treatments independent of the realized demand, with one exception. In the SOFT_HIGH treatment, subjects should not buy any grain when the demand is low, since the predicted market price in this case is only 104 (well below the reserve price of 128).

We find that, as expected, in all treatments and all realized demand, with the exception of SOFT_HIGH treatment when demand is low, subjects buy all eight units—the mean number of units purchased is 8 with a standard error of zero. In the SOFT_HIGH treatment with low demand, subjects buy 4.67 units of grain (with a standard error of 1.05). While subjects buy significantly fewer units compared to other treatments, they also buy significantly higher units than zero (all comparisons are significant at the 95% confidence level).

Figure C.1 shows the total number of units purchased in the auction over rounds for each session of the SOFT_HIGH treatment in Part 2. Note that only the rounds with low demand shocks are shown in this figure, since when the demand is high, individuals always buy eight units (as predicted). Note that buying these units costs subjects the reserve price of 128, which is higher than the market price (both predicted and observed–see Table C.3). Subjects who buy these units are, therefore, losing money. We see from Figure C.1 that, over the rounds, individuals decrease their total purchases in the auction when the demand is low. Since they only face low demand either two or three times in Part 2, it is reasonable to suppose that there were insufficient rounds to allow the market to converge to the theoretically predicted level. While the behavior in the auction is moving towards the predicted levels, these early mistakes might be causing lower carryovers and prices than predicted by our model.

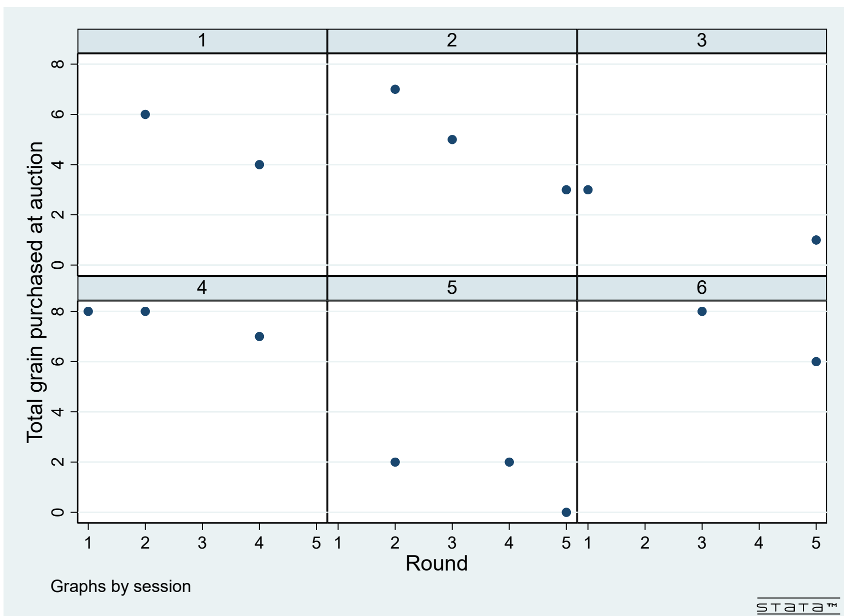

FIGURE C.1. Total grain purchased in SOFT_HIGH for each session when demand is low

Next we look at the prices observed in the auction in each part of the experiment (see Table C.2). For each part, since predictions differ not only based on treatment but also on the realized demand, we calculate the average price for each treatment and demand. Once again, standard errors are based on six independent observations (averages over rounds for a given session, part and demand). Predicted auction prices are based on the predicted spot market prices assuming all subjects bid their true valuations. This is just to create a benchmark. As explained in Section 2, none of our main results depends on the bids submitted in the auction in the baseline and hard floor treatments and, as long as the predicted net additions occur, in the soft floor treatments. As explained above, net additions in the soft floor treatments are exactly what our model predicts with one exception: when floor is high and demand is low. Therefore, even though our model predicted in that case the auction to shut down, there were bids, and the auction price was at the soft floor of 128, as shown in Table C.2. Other than that, Table C.2 shows that all auction prices are substantially lower than the benchmark levels. Nevertheless, there is once again learning. For the BASELINE treatment, where there is no policy change, we see that auction prices are closer to the predicted levels in Part 2 relative to Part 1 for both the low and high demand. The same is true for both the HARD_HIGH and HARD_LOW treatments as well, where the predicted prices do not change between parts.

	Part 1				Part 2			
	Low		High		Low		High	
	Observed	Predicted	Observed	Predicted	Observed	Predicted	Observed	Predicted
BASELINE	69.44 (10.20)	92	145.42 (16.58)	212	77.92 (5.95)	92	164.61 (15.49)	212
SOFT_HIGH	61.19 (3.77)	92	134.78 (12.59)	212	128.00 (0.00)	-	153.92 (6.58)	208
HARD_HIGH	72.86 (6.90)	92	116.53 (11.47)	212	103.67 (8.03)	128	137.39 (13.30)	200
SOFT_LOW	60.00 (5.05)	92	102.94 (8.83)	212	74.58 (2.00)	92	118.78 (16.02)	212
HARD_LOW	61.44 (6.31)	92	124.39 (16.57)	212	73.31 (4.86)	92	138.25 (16.13)	212
SOFT_HIGH_LB	79.22 (6.73)	92	121.47 (16.81)	212	92.97 (6.86)	104	139.39 (18.56)	208
Standard errors in parentheses. There are 6 observations per cell.								

Table C.2. Mean Auction Prices

Finally we also look at the Period 2 prices (see Table C.3). We see that soft price floor increases the Year 2 prices when demand is low but, as expected, the increase is not as large as the hard price floor. This is consistent with the incentive to carry over significantly more when price floor is hard compared to when it is soft.

	Part 1				Part 2			
	Low		High		Low		High	
	Observed	Predicted	Observed	Predicted	Observed	Predicted	Observed	Predicted
BASELINE	99.22 (1.29)	92	218.11 (1.91)	212	97.78 (1.78)	92	216.56 (2.07)	212
SOFT_HIGH	99.33 (1.90)	92	215.89 (2.34)	212	106.78 (1.96)	104	215.39 (1.61)	208
HARD_HIGH	97.06 (2.18)	92	216.28 (2.13)	212	128.00 (0.00)	128	210.5 (1.89)	200
SOFT_LOW	101.11 (1.42)	92	219.33 (1.88)	212	100.06 (1.16)	92	216.33 (2.08)	212
HARD_LOW	98.72 (1.21)	92	217.44 (1.34)	212	102.44 (1.92)	92	218.78 (2.08)	212
SOFT_HIGH_LB	98.56 (1.86)	92	220.56 (1.01)	212	119.11 (1.38)	104	216.39 (1.11)	208
Standard errors in parentheses. There are 6 observations per cell.								

Table C.3. Mean Period 2 Prices

Online Appendix D: Instructions to Subjects in our Laboratory Experiment

In this experiment, you will be linked, via computer, to nine other individuals in this room. Throughout you will interact only with each other, and will never interact with anyone else.

The experiment consists of two parts. In both parts, there will be 5 rounds, 10 rounds in total. During each round, you will have an opportunity to earn experimental dollars (E$). Upon completion of the experiment, one of the 5 rounds from Part 1 and one of the 5 rounds from Part 2 will be randomly chosen. Your earnings in experimental dollars (E$) in these two rounds will be added up and converted to US Dollars ($) at the following rate: 100 Experimental Dollars (E$) = 1.00 US Dollar ($). Your $ earnings (plus the $7 show-up reward) will be paid to you in private at the end of the experiment.

Throughout this experiment, each of you will have the role of a **trader** in a market for a generic commodity ("grain"). You will be asked to decide how much "grain" to buy, sell, or store. Below, we explain what happens in Part 1. There may or may not be some change in Part 2 in how the market works. If there is any change, you will be given a new set of instructions explaining it before Part 2 begins.

Please do not communicate with the other participants during the experiment. Should you have any questions, please raise your hand.

Instructions for Part I:

Product market in grain: Recall that in Part 1, there will be 5 rounds. Each round consists of two "years." In every round, at the beginning of the first year, you will receive both a stock of grain and also a deposit in your money account of E$250. These "starter stocks" of grain and experimental dollars are intended to launch you and will not be repeated in the second year. In the second year, however, you will have the opportunity to buy additional grain at an auction. You will also have the opportunity to store some of the grain from the first year for use in the second year, to supplement any grain you may acquire in the second-year auction. In each year, you will have an opportunity to sell some or all of the grain you own to pre-programmed consumers. We will explain their buying rules below. Any grain you do not sell in the first year is automatically carried into the second year.

Earnings (E$) in each round: Your earnings in each round equals the E$250 of starter funds plus your earnings (E$) from sales to computerized consumers in the two years minus your costs (E$) of buying additional grain at the second year auction. Any grain that you hold at the end of the second year is worthless.

Stock of grain: In every round, each of you will start the first year with 4 units of grain. What you do not sell to the pre-programmed consumers in year one will be automatically stored for your use in year two.

Buying from the grain auction: At the start of the second year, there will be a grain auction where exactly 8 units of grain will be auctioned. You do not have to bid anything. If you want to acquire more grain than you stored, however, you may submit up to two bids. Your bid specifies the most you are willing to pay to buy a unit of grain. You may bid a higher amount (E$) for one unit than for the other. Unsold grain will be retired and will no longer be available to the market. Each of you will bid anonymously without knowing the bids of the other buyers.

Price and acquisitions in the grain auction: All successful bidders pay the same price.

Case 1: If there are 9 or more bids for the 8 units auctioned, then all 8 units are sold. The highest 8 bids will each win one unit of grain. To determine the price every winning bidder pays, we will sort the bids from highest to lowest. The **price** that all successful bidders pay is equal to the value of the highest rejected bid (the bid that is 9th from the highest). If the 8th highest bid is submitted by several bidders, the one who gets the

8th unit auctioned will be decided using a randomizing device giving each tied bidder an equal chance. The price paid in that case is the common bid submitted by these bidders.

Case 2: If there are 8 or fewer bids, then each bidder will win and will pay $E0 (the winners pay nothing!). Any unsold units will be retired from the market.

Selling grain in the market: In each year, you will have the opportunity to sell units of grain to computerized consumers. As a **seller**, you will be asked to post offers for the units of grain that you would like to sell. Your **offer** for a particular unit of grain is the lowest price you would accept for selling the grain. For each of your units of grain, you may post a different offer. Once the grain is sold to computerized consumers, it is no longer available. Any grain you do not sell in the first year is automatically carried into the second year, but any grain that you do not sell in the second year is worthless.

Valuations of computerized (pre-programmed) consumers: There are 48 computerized consumers in the grain market. Each consumer would like to consume one unit of grain in each year but values a unit consumed in year 2 differently than one consumed in year 1. **Table 1** (Table D.1. in this Online Appendix) summarizes, in its 48 rows, the most each of these 48 consumers would be willing to pay (their "valuations") for one unit of grain consumed in year 1 or year 2. We have ranked the 48 consumers from highest to lowest according to their valuations. The consumer with the highest valuation receives the index 1, the consumer with the lowest valuation receives the index 48. Column 1 of Table 1 lists the indices of the consumers. Column 2 indicates the valuations of the consumers in year 1. In year 2, consumer willingness to pay for grain can be either high or low. Column 3 indicates the valuations of the consumers in year 2 when demand is high, that is if every consumer has a higher valuation than in year 1. Column 4 indicates the valuations of consumers if demand is low, that is if every consumer has a lower valuation than in year 1. High and low demand occur with equal chances and you will be told at the start of year 2 which of these two situations has occurred.

The highest valuation for one unit is E$198 in year 1. In year 2, it is E$258 if demand is high and E$138 if demand is low. Note that as one moves down one row, the valuations in Columns 2, 3, and 4 decrease by E$2.

Market price for grain: In the market some of the traders (including you) will succeed in selling units to some of the computerized consumers. All transactions will take place at the same price.

In order to determine who makes a trade, we use the following procedure. First, we rank offers of the traders from lowest to highest, giving any tied offers distinct ranks, and we rank valuations of the computerized consumers from highest to lowest. Second, we pair the lowest offer with the highest valuation, the second lowest offer with the second highest valuation, until we run out of offers. Third, we discard incompatible pairs—pairs where the buyer values the unit at less than the seller requires to sell it. Every undiscarded pair results in a transaction. All transactions take place at the same price: the lowest undiscarded valuation. We refer to this price as the "market price."

Example: We can use the example in **Table 2** (Table D.2 in this Online Appendix) to describe how we determine the market price and quantity. Suppose it is year 2 and the demand is high. The valuations of the 48 consumers are ranked highest to lowest in Column 2 (for your convenience, we have duplicated the information in Column 3 of the previous table) and the offers of the traders are ranked lowest to highest in Column 3. **Find the lowest row such that the valuation in Column 2 is greater than or equal to the offer in Column 3. The entry in Column 1 is the quantity sold and the entry in Column 2 is the price per unit (E$) that every trader receives.** So in this example 42 units would be sold at a price of E$176. At that price, there are exactly 42 buyers willing to pay E$176 and 42 units offered by traders at a price E$176.

Why not some other row? Try it. A row further toward the bottom of the table results in a price so low that there would be more buyers willing to purchase at that price than units offered at that price. Similarly a row closer to the top of the table results in a price so high that there would be fewer buyers willing to purchase at that price than units offered at that price.

Note that the price earned from a sale in the grain market can be greater than or less than the price paid to acquire a unit in the auction.

Summary of Part 1:

There will be 5 rounds of two-year grain market games. At the start of each two-year round, you will receive (1) E$250 and (2) 4 units of grain. At the start of second year, you will also have the opportunity to buy additional grain in an auction where all winning bidders pay the same price.

In both years, you will make offers to sell grain to pre-programmed consumers. In the first year, any grain you do not sell will be carried automatically into the second year, and will supplement any grain you may acquire in the second-year auction. At the end of the second year, any unsold grain will have no value.

All grain sold to pre-programmed consumers will be sold at a single market price that is set so that no unit is sold unless the most consumers are willing to pay for it exceeds (or equals) the asking price (offer) of the trader who owns it. As for the traders who make no sales and the consumers who make no purchases, even the consumer with the highest valuation among them is unwilling to pay as much as the trader with the smallest asking price requires.

Instructions for Part 2 (Baseline Treatment)

Part 2 will consist of another 5 rounds of the same two-year grain market games as in Part 1. There are NO changes in the rules regarding how the markets operate.

Instructions for Part 2 (Reserve Price Treatment–SOFT_HIGH)

Part 2 consists of 5 rounds of two-year grain markets. There is only one change: now no bids lower than E$128 may be submitted in the grain auction. When Part 2 starts, this minimum allowable bid (E$128) will appear on everyone's computer screen.

Price in the grain auction: The price is calculated in the same way as in Part 1, with only one exception: *If the number of bids is less than or equal to 8 then every bid will be successful and will pay E$128 on each winning bid. Therefore, the price in the grain auction will never fall below E$128.*

Note that, despite the requirement that bids in the auction equal or exceed E$128, the price in the grain market can still fall below E$128.

Instructions for Part 2 (Reserve Price Treatment–SOFT_LOW)

Same as the instructions for SOFT_HIGH with the exception that 128 is replaced with 68.

Instructions for Part 2 (Hard Price Floor Treatment–HARD_HIGH)

Part 2 consists of 5 rounds of two-year grain market games. There is only one change: in Year 2, there is now a special buyer in the grain market in addition to the pre-programmed consumers of Part 1 who is willing to buy unlimited amounts at a price of E$128. This changes the bid schedule in the following way: all bids lower than E$128 are replaced by E$128 in Column 4. This is shown in Table 3 (Table D.3 in this Online Appendix). Note that, because of this change, the price in the grain market will never fall below E$128 in Year 2.

Other than the change in the bid schedule in the second year, two-year markets operate exactly the same way as in Part 1.

Instructions for Part 2 (Hard Price Floor Treatment–HARD_LOW)

Same as the instructions for HARD_HIGH with the exception that 128 is replaced with 68 and Table 3 is adjusted accordingly.

Instructions for Part 2 (Follow-up Experiment–SOFT_HIGH_LB)

Part 2 consists of 5 rounds of two-year grain market games. There is only one change: in Year 2, there is now a special pre-programmed buyer in the grain market, in addition to the pre-programmed consumers of Part 1, who is willing to buy up to eight units of grain at a price of E$128. Hence, if demand is LOW, once the valuations reach 128, they no longer descend further until the special pre-programmed buyer has purchased an 8th unit. This changes the "Year 2 Valuations (E$) if Demand is LOW" in the following way: E$128 now repeats nine times in Column 4. This is shown in Table 3 (Table D.4 in this Online Appendix). Note that, despite this change, the price in the grain market may fall below E$128 in year 2.

Other than the change in the bid schedule in the second year, two-year markets operate exactly the same way as in Part 1.

(1)	(2) Year 1 Valuations (E$)	(3) Year 2 Valuations (E$) if Demand is HIGH	(4) Year 2 Valuations (E$) if Demand is LOW
1	198	258	138
2	196	256	136
3	194	254	134
4	192	252	132
5	190	250	130
6	188	248	128
7	186	246	126
8	184	244	124
9	182	242	122
10	180	240	120
11	178	238	118
12	176	236	116
13	174	234	114
14	172	232	112
15	170	230	110
16	168	228	108
17	166	226	106
18	164	224	104
19	162	222	102
20	160	220	100
21	158	218	98
22	156	216	96
23	154	214	94
24	152	212	92
25	150	210	90
26	148	208	88
27	146	206	86
28	144	204	84
29	142	202	82
30	140	200	80
31	138	198	78
32	136	196	76
33	134	194	74
34	132	192	72
35	130	190	70
36	128	188	68
37	126	186	66
38	124	184	64
39	122	182	62
40	120	180	60
41	118	178	58
42	116	176	56
43	114	174	54
44	112	172	52
45	110	170	50
46	108	168	48
47	106	166	46
48	104	164	44

Table D.1. Valuation Schedule (in E$) of the Pre-Programmed Buyers in the Grain Market

(1)	(2) Year 2 Valuations of Pre-Programmed Consumers if Demand is HIGH	(3) Offers of Traders
1	258	50
2	256	50
3	254	65
4	252	70
5	250	72
6	248	80
7	246	85
8	244	87
9	242	90
10	240	91
11	238	99
12	236	100
13	234	100
14	232	100
15	230	100
16	228	101
17	226	105
18	224	105
19	222	105
20	220	105
21	218	105
22	216	105
23	214	108
24	212	109
25	210	110
26	208	110
27	206	110
28	204	110
29	202	115
30	200	120
31	198	120
32	196	130
33	194	135
34	192	135
35	190	140
36	188	142
37	186	150
38	184	150
39	182	150
40	180	152
41	178	160
42	176	170
43	174	175
44	172	180
45	170	180
46	168	185
47	166	
48	164	

TABLE D.2. How the Price is Calculated in the Grain Market

(1)	(2) Year 1 Valuations (E$)	(3) Year 2 Valuations (E$) if Demand is HIGH	(4) Year 2 Valuations (E$) if Demand is LOW
1	198	258	138
2	196	256	136
3	194	254	134
4	192	252	132
5	190	250	130
6	188	248	128
7	186	246	128
8	184	244	128
9	182	242	128
10	180	240	128
11	178	238	128
12	176	236	128
13	174	234	128
14	172	232	128
15	170	230	128
16	168	228	128
17	166	226	128
18	164	224	128
19	162	222	128
20	160	220	128
21	158	218	128
22	156	216	128
23	154	214	128
24	152	212	128
25	150	210	128
26	148	208	128
27	146	206	128
28	144	204	128
29	142	202	128
30	140	200	128
31	138	198	128
32	136	196	128
33	134	194	128
34	132	192	128
35	130	190	128
36	128	188	128
37	126	186	128
38	124	184	128
39	122	182	128
40	120	180	128
41	118	178	128
42	116	176	128
43	114	174	128
44	112	172	128
45	110	170	128
46	108	168	128
47	106	166	128
48	104	164	128

TABLE D.3. Valuation Schedule of the Pre-Programmed Buyers [Hard Floor Treatment]

(1)	(2)	(3)	(4)
	Year 1 Valuations (E$)	**Year 2 Valuations (E$) if Demand is HIGH**	**Year 2 Valuations (E$) if Demand is LOW**
1	198	258	138
2	196	256	136
3	194	254	134
4	192	252	132
5	190	250	130
6	188	248	128
7	186	246	128
8	184	244	128
9	182	242	128
10	180	240	128
11	178	238	128
12	176	236	128
13	174	234	128
14	172	232	128
15	170	230	126
16	168	228	124
17	166	226	122
18	164	224	120
19	162	222	118
20	160	220	116
21	158	218	114
22	156	216	112
23	154	214	110
24	152	212	108
25	150	210	106
26	148	208	104
27	146	206	102
28	144	204	100
29	142	202	98
30	140	200	96
31	138	198	94
32	136	196	92
33	134	194	90
34	132	192	88
35	130	190	86
36	128	188	84
37	126	186	82
38	124	184	80
39	122	182	78
40	120	180	76
41	118	178	74
42	116	176	72
43	114	174	70
44	112	172	68
45	110	170	66
46	108	168	64
47	106	166	62
48	104	164	60

Table D.4. Valuation Schedule of the Pre-Programmed Buyers [Soft Floor Treatment with Limited Buyback]