European Economic Review 159 (2023) 104583

Contents lists available at ScienceDirect

European Economic Review

journal homepage: www.elsevier.com/locate/eer

The effects of seemingly nonbinding price floors: An experimental analysis[☆]

Stephen Salant[a], William Shobe[b], Neslihan Uler[c,*]

[a] University of Michigan, United States of America
[b] University of Virginia, United States of America
[c] University of Maryland and University of Michigan, United States of America

ARTICLE INFO

JEL classification:
C6
C92
D0
Q1
Q5

Keywords:
Price floors
Emissions markets
Agricultural price supports
Commodity markets
Laboratory experiments

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.

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

☆ 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 Department of Agricultural and Resource Economics at the 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.

* Corresponding author.
E-mail addresses: ssalant@umich.edu (S. Salant), shobe@virginia.edu (W. Shobe), neslihan@umd.edu (N. Uler).

https://doi.org/10.1016/j.euroecorev.2023.104583
Received 2 October 2022; Received in revised form 28 August 2023; Accepted 1 September 2023
Available online 9 September 2023

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

Our companion paper (Salant et al., 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 whether individuals react to seemingly nonbinding price floors in a way predicted by theory and whether action at a distance is stronger with a hard floor than with a soft floor. To answer these questions, 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.

Switching from the infinite-horizon formulation of our companion paper (Salant et al., 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 et al. (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, Lee (1978) showed that a hard price ceiling on exhaustible resources such as oil will affect today's price, even if the price

[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.

[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 collars. 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.

[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.

would not hit the ceiling until well into the future. The "action-at-a-distance" hypothesis 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 et al. (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 et al. (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 et al. (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 et al. (2016), and Friesen et al. (2022). Multiperiod experiments involving price controls like our own began with Stranlund et al. (2014), who explore the effects of banking and price controls on price volatility. Cason et al. (2023) 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 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 et al., 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 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.

[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).

[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.

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

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:

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

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 $f \le 92,\ p_h^H = 212, p_l^H = 92$, 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 $f \le 92$, 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 $g_l^H \ge 0$. 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 second period price can fall below the floor in the low-demand state. Note that, in theory, the soft floor can equally well be achieved with a limited buyback. We summarize these results in Fig. 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.

[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.

[15] In this example, the equilibrium price with no floor is 152. But as Fig. 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, if set above it, would raise the wage only to the level of the floor, never above it.

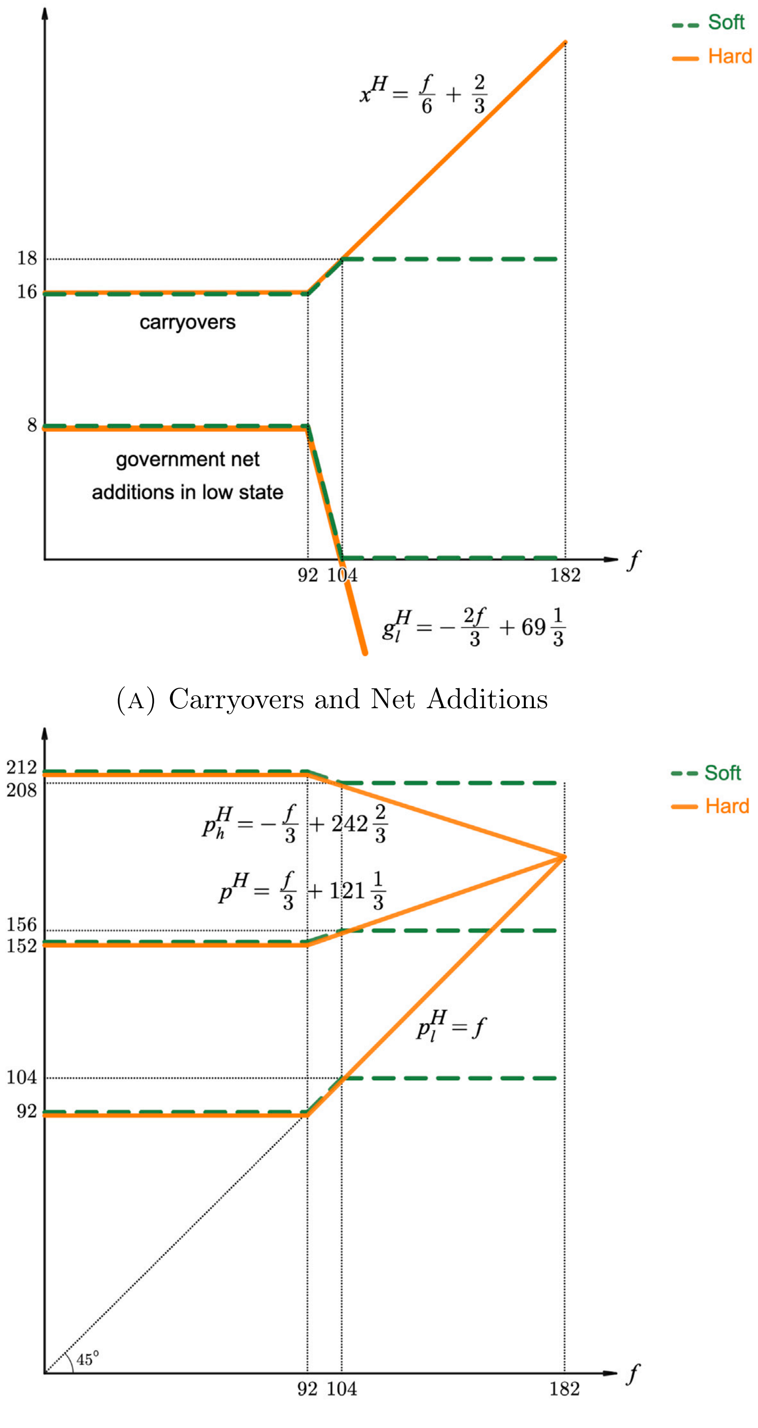

(B) Prices

Fig. 1. Graphical representation of the model for the parameters used in our experiment.

Table 1
Predictions in a two-period example.

	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 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 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.

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 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

[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).

[17] As can be seen in Fig. 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.

[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".

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 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.

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. 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.

[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.

[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).

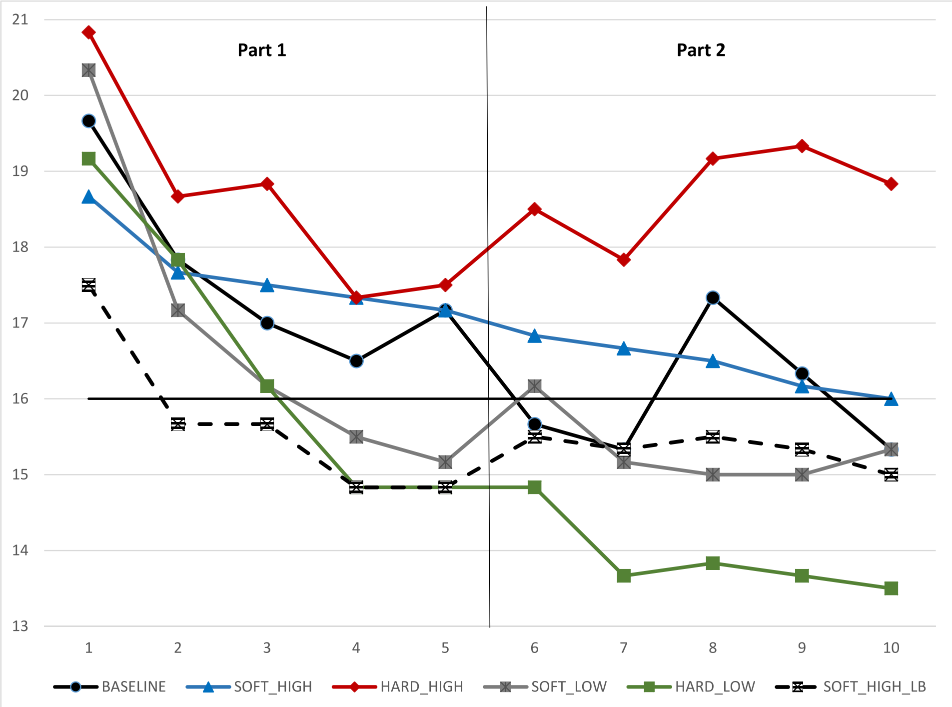

Fig. 2. Average carryover, all rounds (5 in part 1 and 5 in part 2).

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.

Fig. 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. Fig. 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 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 HARD_HIGH treatment are one-sided. In that case, naturally the significance level of these results would be higher. For example,

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

Table 2
Mean carryover.

	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 3
Are carryovers different between treatments in Part 2?

	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.

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, Fig. 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 Fig. 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.

Table 4
OLS regression analysis.

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.

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

[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.

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 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 establishes that action at a distance, although predicted by conventional theory, does not occur in our high soft-floor case when it is implemented as a limited buyback. Since it does occur in our high hard-floor case which uses the same format (and since it does occur in Holt-Shobe's soft-floor implementation despite its greater complexity), it seems to us unlikely that the absence of action at a distance is attributable to the complexity of the experimental environment.[25] Only further research can determine the actual source of this departure from the prediction of standard theory.

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.[26] The mechanism we identify likely explains the results observed by Holt and Shobe (2016) in their laboratory experiments on EU-ETS Market Stability Reserve.

[25] Having said that, we acknowledge that we cannot completely eliminate that possibility, and, as pointed out to us by an anonymous referee, experts in the field might be able to figure out optimal behavior while our student subjects might fail to do so in a laboratory experiment.

[26] Our analysis justifies the conjecture in the *Financial Times* Terazona (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.

Using linear demands, we derived closed-form results more suitable for testing than the infinite-horizon formulation of Salant et al. (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 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.[27] This result is inconsistent with the pattern observed in Holt and Shobe (2016). This unexpected result will be a fruitful area for future research.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.euroecorev.2023.104583.

References

Burtraw, Dallas, Palmer, Karen, Kahn, Danny, 2010. A symmetric safety valve. Energy Policy 38, 4921–4932.
Cason, Timothy N., Stranlund, John K., de Vries, Frans P., 2023. Investment incentives in tradable emissions markets with price floors. J. Assoc. Environ. Resour. Econ. 10 (2), 283–314.
Cottle, Rex L., Wallace, Myles S., 1983. Economic effects of non-binding price constraints. J. Ind. Econ. 31 (4), 469–474.
Deaton, Angus, Laroque, Guy, 1992. On the behaviour of commodity prices. Proc. Natl. Acad. Sci. 59, 1–23.
Dow, James, Olson, Lars, 2001. An analytic solution to a stochastic consumption/saving problem with liquidity constraints. South. Econ. J. 459–464.
Fell, Harrison, Burtraw, Dallas, Morgenstern, Richard D., Palmer, Karen, 2012. Soft and hard price collars in a cap-and-trade system: A comparative analysis. J. Environ. Econ. Manag. 64 (2), 183–198.
Fell, Harrison, Morgenstern, Richard D., 2010. Alternative approaches to cost containment in a cap-and-trade system. Environ. Resour. Econ. 47 (2), 275–297.
Fischbacher, Urs, 2007. z-Tree: Zurich toolbox for ready-made economic experiments. Exp. Econ. 10 (2), 171–178.
Friesen, Lana, Gangadharan, Lata, Khezr, Peyman, MacKenzie, Ian A., 2022. Mind your Ps and Qs! variable allowance supply in the US regional greenhouse gas initiative. J. Environ. Econ. Manag. 112, 102620.
Gardner, Bruce, 1979. Optimal Stockpiling of Grain. Lexington Books.
Greiner, Ben, 2015. Subject pool recruitment procedures: Organizing experiments with ORSEE. J. Econ. Sci. Assoc. 1, 114–125.
Gustafson, Robert L., 1958. Carry-Over Levels for Grains. Tech. Bulletin 1178, Department of Agriculture, Washington.
Holt, Charles A., Shobe, William M., 2016. Price and quantity collars for stabilizing emission allowance prices: Laboratory experiments on the EU ETS market stability reserve. J. Environ. Econ. Manag. 76, 32–50.
Isaac, R. Mark, Plott, Charles R., 1981. Price controls and the behavior of auction markets: An experimental examination. Amer. Econ. Rev. 71, 448–459.
Kohn, Meir, 1978. Competitive speculation. Econometrica 46, 1061–1076.
Krugman, Paul, 1991. Target zones and exchange rate dynamics. Q. J. Econ. 106, 669–682.
Lee, Dwight R., 1978. Price controls, binding constraints, and intertemporal economic decision making. J. Political Econ. 86 (2), 293–301.
MacKenzie, Ian A., 2022. The evolution of pollution auctions. Rev. Environ. Econ. Policy 16 (1), 1–24.
Newbery, David M.G., Stiglitz, Joseph E., 1981. The Theory of Commodity Price Stabilization. Oxford University Press.
Olson, Lars J., 1989. Stochastic growth with irreversible investment. J. Econom. Theory 47, 101–129.
Perino, Grischa, Willner, Maximilian, Quemin, Simon, Pahle, Michael, 2022. The european union emissions trading system market stability reserve: Does it stabilize or destabilize the market? Rev. Environ. Econ. Policy 16 (2), 338–345.
Perkis, David F., Cason, Timothy N., Tyner, Wallace E., 2016. An experimental investigation of hard and soft price ceilings in emissions permit markets. Environ. Resour. Econ. 63 (4), 703–718.
Salant, Stephen W., 1983. The vulnerability of price stabilization schemes to speculative attack. J. Polit. Econ. 91 (1), 1–38.
Salant, Stephen, Shobe, William, Uler, Neslihan, 2022. The effects of nonbinding price floors. Eur. Econ. Rev. 145, 104122.
Samuelson, Paul, 1971. Stochastic speculative price. Proc. Natl. Acad. Sci. 68, 335–337.
Scheinkman, Jose A., Schechtman, Jack, 1983. A simple competitive model with production and storage. Rev. Econom. Stud. 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. J. Environ. Econ. Manag. 40, 189–210.
Smith, Vernon L., Williams, Arlington W., 1981. On nonbinding price controls in a competitive market. Amer. Econ. Rev. 71 (3), 467–474.
Solow, Robert, 1974. The economics of resources or the resources of economics. Amer. Econ. Rev. 1–14.
Stranlund, John K., Murphy, James J., Spraggon, John M., 2014. Price controls and banking in emissions trading: An experimental evaluation. J. Environ. Econ. Manag. 68, 71–86.
Svensson, Lars E.O., 1992. An interpretation of recent research on exchange rate target zones. J. Econ. Perspect. 6 (4), 119–144.
Terazona, Emiko, 2013. Beijing's rice subsidies buoy imports. Financial Times.
Wright, Brian D., Williams, Jeffrey C., 1982. The economic role of commodity storage. Econom. J. 92 (367), 596–614.

[27] We do not conclude that soft floors never work. Our paper suggests the existence of an environment in which our theory predicts action at a distance and yet surprisingly we do not see it. This raises the possibility that, under some circumstances outside the laboratory, soft floors might not work as effectively as theory predicts.