The Effects of "Non-binding" Price Floors

by Stephen Salant, William Shobe, and Neslihan Uler

Abstract

Hard price floors are a standard form of agricultural price support while soft floors (auction reserve prices) are often included in cap-and-trade programs. A hard floor is a government commitment to purchase at the floor price, in unlimited amounts, whatever commodity is offered for sale, whereas a soft floor can be regarded as a commitment to buy at that price up to the limited amount of commodity the government auctions each period. We investigate theoretically the effect of each type of floor in the stochastic, competitive storage model adopted independently by agricultural and environmental economists. We find that, even if the floor is inserted below the initial spot price, that price will, under specified conditions, jump up—a phenomenon that we call "action at a distance." Moreover, the jump is larger with a hard floor than with a soft floor. We distinguish our results from those of Krugman (1991) where a foreign exchange rate is supported by a hard floor. Unlike his case, the fundamental in our model does not execute a random walk and continues to strictly affect the spot price even in the neighborhood of the floor: there is no "smooth pasting" result. We are testing our theoretical predictions using laboratory experiments.

Preliminary draft: Please do not quote without permission

The Effects of Non-binding" Price Floors

Stephen Salant, William Shobe and Neslihan Uler*

October 31, 2018

1 Introduction

In classes on Economics Principles where the competitive equilibrium in a market is depicted by the intersection of a supply and demand curve, a price floor can displace this equilibrium but only if it is inserted above the equilibrium price. A price floor inserted below it is "non-binding or, as *Wikipedia* asserts at the outset of its article on Price Floors, ineffective.

Static diagrams showing that the policy is non-binding are very misleading if the "good" depicted is storable as is the case of grains (corn, wheat, rice, etc.), extracted resources (oil, metals, etc.), emissions permits (for CO_2, SO_2), and foreign exchange. When the asset is storable, a policy that affects only its future price will typically affect its current price.

In the case of permits allowing carbon emissions, variable weather induces *demand* shocks. Variation in heating or cooling results in variation in carbon emissions and, under cap-and-trade programs, the government requires each regulated firm to cover emissions during the period by surrendering permits.[1] The government can distribute allowances for free, sell them at auction, or some combination of the two.

When permit prices are low, regulated firms find it cheaper to surrender permits than to incur the cost of reducing emissions (abating); when permit prices are high, however, it becomes cheaper to incur abatement costs instead and firms surrender fewer permits. Hence, the demand for permits to surrender at time t depends on the current permit price and on the realized demand shock.

Any permits not surrendered in the current period are carried forward into the next period. Cap-and-trade programs typically include auctions so that the permits that were surrendered in the last period are at least partially replenished. Assume that g permits are auctioned in each period.

*Salant: University of Maryland, salant@umich.edu. Shobe: University of Virginia, shobe@virginia.edu. Uler: University of Maryland, neslihan@umd.edu. The authors would like to thank Dallas Burtraw, Julien Daubanes, Harrison Fell, Gérard Gaudet, Chris House, Christophe Gouel and Kyle Meng for their valuable comments on earlier drafts. We thank David Borszich, Ia Vardishvili, and Haozhu Wang for their skillful research assistance. Salant wishes to acknowledge useful discussions with Yichuan Wang and Ping Han, two undergraduates in his research seminar with whom he worked at the inception of this research project. He also wishes to thank the Erb Institute at University of Michigan and AREC at University of Maryland for financial support of this research. We gratefully acknowledge financial support from the International Foundation for Research in Experimental Economics (IFREE). We thank John Jensenius for his help with programming and running the experiment. We also thank Andrew Card and Aldo Gutierrez Mendieta for superb research assistance.

[1] We conduct our analysis in discrete time and define one period as the length of time between true-ups.

In the case of grains, harvests are subject to weather-induced *supply* shocks. Any grain not sold to consumers or grain processors in period t is carried forward into the next period. The current harvest at least partially replenishes the grain that was sold to consumers and grain processors in the previous period.

Weather fluctuations and other shocks can cause wide and unpredictable swings in prices in these markets. To prevent prices from falling too far, governments often intervene to provide a price floor. In permit markets, governments may reduce the amount auctioned when the market price is too low. They do so by imposing a minimum or reserve price (f) on bids in the auction (a so-called "soft floor). Bids below the reserve price at the auction are disqualified. Since there is no reason to bid more than the current market price and competition among bidders keeps the winning bids from falling much below the market price, the auction price tends to be equal to or just below the market price. If the price is above the soft floor, all g permits are auctioned. If the price is below the floor, no permits are auctioned. If the price is exactly equal to the floor, g or fewer permits are auctioned. Since the government can do no more using a reserve price than to shut down the current auction, the market price can fall below a soft floor.

In grain markets, prices are prevented from falling too far by means of a buyback (a so-called "hard floor"). The government stands ready to purchase as much as farmers want to sell at the floor price (f). Since the government is always willing to buy at the floor, the market price never falls below it.

Although these three policies (no floor, a soft floor and a hard floor) seem qualitatively dissimilar, they can be viewed within a unified framework. Each can be regarded as an auction of g permits with no reserve price followed by an anticipated constrained buyback at exogenous price f. The buyback is constrained to zero in the case of no floor, to g in the case of a soft floor, and is unconstrained in the case of a hard floor.

For decades, economists have pondered the consequences of imposing such floors on the price of the storable asset. As we will show, inserting either type of floor under what the the spot price would have been in the absence of a floor may cause that spot price to jump up, a phenomenon we refer to as "action at a distance." When a soft floor causes the price to jump, a hard floor at the same level causes the price to jump at least as high. We derive these predictions theoretically, illustrate them by simulations and test them in the laboratory. We conclude this introduction by reviewing the relevant articles in the agricultural price-support literature and the emissions trading literature to distinguish our contribution from what has preceded it.

For more than half a century, hard price floors have been used in agricultural programs to support grain prices. In recent decades, soft floors have been included in cap-and-trade programs to cushion drops in permit prices. In each case, a literature developed to understand the consequences of these policies. The literature on agricultural price supports developed in the 1950's and the literature on emissions permits in recent decades. Both literatures study the same problem. Although the price of foreign exchange is often also supported by hard floors, the assets we study differ from foreign exchange: they lack any "convenience yield" and must be continually replenished by har-

vests (or auctions) to offset their disappearance through consumption (or true-ups). Nonetheless, a comparison to Krugman's (1991) analysis of the effect of target zones on foreign exchange rates is illuminating, and we conduct it after presenting our model.

Both the literature on agricultural price supports and the literature on emissions permits analyze the stochastic, competitive equilibrium that arises when profit-maximizing private agents can sell an asset (be it a bushel of wheat or a permit to emit a ton of carbon) at the current price or store it for sale in the future. Both literatures recognize that unpredictable weather and other shocks continually disturb these markets.

Eventually, economists studying the prices of emissions permits recognized that the model they were developing bore a striking resemblance to the earlier stochastic agricultural storage model.[2] In the following subsections, we review what these two literatures concluded about the consequences of price floors in order to distinguish our contribution from what preceded it. We begin with the older literature.

1.1 Hard Price Floors in Agricultural Commodity Agreements

The early literature, (Working, 1953) and (Johnson, 1954), expressed the view that the unanticipated imposition of a hard floor strictly below the current price may cause the current price to jump up. For example, Holbrook Working noted that the price support program for wheat "will give growers sound reason to expect prices in its earlier years to average *above* an equilibrium level." Similarly, D. Gale Johnson noted that "the existence of price supports at relatively profitable levels acts as an inducement to produce *even when market prices have been above support levels*, as has been true much of the time since 1941" [our emphasis]. However, clarification of the *mechanism* responsible for such "action at a distance" and the specification of sufficient conditions for it to occur had to await the development of intertemporal, stochastic models of grain carryovers.[3]

Using dynamic programming, Gustafson (1958) approximated the optimal storage of grain numerically when demand is linear and the harvest is uncertain. However, the nonnegativity condition on stocks prevented him from deriving the optimal storage rule analytically. Samuelson (1971) showed analytically that, for general utility functions, a unique optimal storage rule exists. Both

[2]Schennach (2000) was among the first to note the close connection between models of commodity markets and the modeling of intertemporal allowance markets under uncertainty. In adding uncertainty to the Rubin intertemporal framework, Schennach noted that the current allowance price depends on the expected discounted price in the future. A positive bank (positive storage of permits) leads to allowance prices expected to rise at the rate of interest. As she noted, however, an unanticipated increase in the possibility that the bank might fall to zero at some point in the future (a "stockout") raises the current price. Schennach noted that the current effect of the possibility of future stockouts had been extensively analyzed in the commodity markets literature.

[3]There are also theoretical and experimental analyses of static models that are somewhat related literature (Eeckhoudt and Hansen (1980), Cottle and Wallace (1983)) on how a competitive firm choosing its output and then drawing a price from an exogenous price distribution would react if the government truncated that price distribution by imposing a hard floor. In an experimental investigation of non-binding price controls in a static setting, Isaac and Plot(1981) were able to reject the hypothesis that non-binding price controls act as focal points for bidders, but they were not able to conclude that non-binding controls near the competitive equilibrium (CE) push prices away from the CE. Smith and Williams (1981) show that non-binding price controls do have an effect in the static setting. The controls, although non-binding, restrict the distributions of bids and offers that participants make in the double auction market setting. This results in the shifting of the CE price away from the non-binding control.

Gustafson and Samuelson appealed to the first welfare theorem to reinterpret the solutions to their optimization problems as competitive equilibria.

Gardner (1979) developed a computerized model to analyze the competitive equilibrium in a stochastic storage model when the government defended a price *band*. To defend the floor, the government purchases anything offered at the floor price; to defend the ceiling, the government sells at the ceiling price anything demanded up to its entire grain stockpile. Since the competitive equilibrium in this case solves no known optimization problem, Gardner abandoned the optimization framework of Gustafson and Samuelson. Instead he approximated the sequence of equilibrium price functions directly.

Salant (1980; 1983) analyzed the competitive equilibrium using the same approach as Gardner but was able to obtain analytical results because he assumed the floor and ceiling coincide (a "price peg" of the kind previously used for gold and foreign exchange).[4] Williams and Wright (1991), using Gardner's code, discussed the positive and normative effects of price bands. In 1988, Miranda and Helmberger (1988) and Wright and Williams (1988) independently analyzed price stabilization schemes in a rational expectations framework. Although one can sometimes infer from their simulations the occurrence of "action at a distance" when a hard floor is imposed, theory is necessary to clarify the underlying mechanism and to identify in general when such a phenomenon would occur.

In our view, the contribution which is closest in spirit to our paper is by Dwight Lee (1978) although his article concerns the effect of a hard *ceiling* which is strictly above the laissez-faire price in a *deterministic* Hotelling model of an *exhaustible* resource. He showed that the unanticipated imposition of a hard ceiling will depress the current price if it is foreseen that the price will eventually reach the hard ceiling. Our analysis shows how this result can be generalized and applied to stochastic settings with price floors or ceilings.

1.2 Price Floors in Emissions Permit Markets

In the absence of uncertainty, Montgomery's (1972) demonstrated a key equivalence between price and quantity instruments for controlling pollution. Weitzman (1974) showed that, under uncertainty over compliance costs, a fixed supply of allowances that is completely unresponsive to price breaks this symmetry between the two types of instruments and can lead to one being preferred over the other in particular cases. Roberts and Spence (1976) proposed a mixed regime where subsidies and penalties at high and low price extremes could limit the range of costs of high price volatility in the allowance market. Refined versions of the Roberts and Spence idea have become known as "price collars" in the emission market context.

A great deal has been written, since the initial contribution by Roberts and Spence (1976), about efforts to rein in price volatility in emission markets. The original concern was that, with fluctuating demand, a vertical supply curve for allowances would induce excess price volatility. The

[4]Newbery and Stiglitz (1981) provide an in depth analysis of commodity price stabilization, including (in their Chapter 29.2.1) a useful summary of the speculative attack result in Salant (1979).

price volatility is induced by the regulator's setting of a fixed cap on emission that does not respond in any way to signals of scarcity or plenty in the allowance market. The volatility could be a short-run phenomenon, say due to weather conditions, business cycles or other demand drivers. Or the movement of price away from its expected range might reflect a longer run shift in relative scarcity of allowances due to technological change, economic growth or regulatory policies. Particular attention has been paid to long-run imbalances that might occur in the effort to control greenhouse gas emissions. In the emission trading literature, the term "safety valve" is often used for efforts to manage the range of price variability.

Price collar proposals all boil down to efforts to add price responsiveness to the allowance supply curve.[5] At some high trigger price, more allowances are made available than would be the case at a price below the trigger (Pizer, 2002; McKibbin and Wilcoxen, 2002). Low price collars modeled in the literature usually take one of two forms, a reserve price in allowance auctions or a government offer to buy back allowances at the trigger price (Fell and Morgenstern, 2010; Fell et al., 2012; Murray, Newell and Pizer, 2009).

Burtraw, Palmer and Kahn (2010) provided a detailed assessment of the arguments for a price collar and clearly enunciated the case for expecting action at a distance. Burtraw et al. noted that "even if the one-sided safety valve never does bind, its introduction to a cap and trade program affects the expectation of future emissions levels and allowance prices and thereby the expectations about the payoff from various investment strategies" (2010, pp. 4922). They used a simulation model to show how a price collar changes the probability distribution of future allowance price outcomes and how this feeds back to current and future paths for electricity market outcomes. They argued that a high price safety valve by itself may not be sufficient, since it truncates the upper end of the price distribution without addressing the empirically more relevant case of unexpectedly low allowance prices.

There have been numerous experiments involving emission markets with features designed to reduce price variability. Most of these price collar experiments implement controls by adjusting the available supply of allowances.[6]

Perkis, Cason and Tyner (2016) used a static experimental setting to compare hard and soft price collars. As expected, they found that, for binding controls, soft collars were not as effective as hard ones. Stranlund, Murphy and Spraggon (2014) used a multi-period setting to show that, even in the presence of banking, price collars may reduce price volatility, since market participants may not be able to smooth compliance costs sufficiently with banking alone and that reduced price volatility may result in increased emission volatility.

In experiments on price collars in the European Emission Trading System (EU ETS), Holt and Shobe (2016) had subjects acquire emission allowances in an auction with both a reserve price and

[5]The European Union Emission Trading System has implemented a "quantity collar" mechanism that uses the size of the bank as a trigger for making temporary adjustments to the number of allowances made available in a given year. This mechanism is not an adjustment to allowance supply and is not considered here (Holt and Shobe, 2016).

[6]One key exception is the EU ETS market stability reserve (MSR), which did not, in its initial implementation, change the aggregate supply of allowances, but rather attempts to limit the quantity of banked allowances. See (Holt and Shobe, 2016)

a high trigger price for the release of additional allowances. Allowances could be banked for future use, inducing temporal smoothing by subjects. Their results suggest that, *even when the reserve price is strictly below the spot price, raising the auction reserve price raised the observed spot price*, consistent with our prediction of "action at a distance" with a soft collar.

While providing suggestive discussions and simulation results, none of the papers in either of these two literatures characterizes analytically sufficient conditions for "action at a distance" to occur, none shows that the effect is stronger for hard floors than for soft floors, and none tests its theoretical results in a laboratory setting.

We proceed as follows. In section 2, we describe the standard infinite-horizon, stochastic storage model in the absence of a price floor and show how to incorporate hard and soft floors into it. In section 3, we explain the price consequences of imposing either of these two policies and compare their effects. We conclude section 3 by comparing our model to Krugman's (1991) target-zone model. In section 4, we present a tractable, two-period model where the mechanism at work in the infinite-horizon case is predicted to generate action at a distance. Section 5 outlines the experimental design we use to test the predictions of this two-period model (we include the instructions as an appendix). Section 6 concludes the paper.

2 Adapting the Standard Storage Model of Gustafson-Samuelson to Include Price Floors

We consider an infinite-horizon, stochastic competitive model of storage where "speculators" are risk neutral and have rational expectations. In period t ($t = 1, \ldots$), agents enter with aggregate stock $s_t \geq 0$, observe the shock realization α_i ($i = 1, \ldots, N$), and decide how much of the stock to sell currently and how much to carry into the next period. The carryover from period t is denoted x_t and becomes the next period stock:

$$s_{t+1} = x_t \text{ for } t = 1, \ldots \tag{1}$$

Shocks are drawn from a stationary probability mass function with Prob $(\alpha = \alpha_i) = \pi_i$, where $\alpha_1 > \ldots > \alpha_N$ and $E(\tilde{\alpha}) = \sum_{i=1}^{N} \pi_i \alpha_i = 0$. The properties of the stochastic competitive equilibrium of the Gustafson-Samuelson model are intuitive. The discounted price expected at t to prevail at $t+1$ can never strictly exceed the price at t; otherwise, risk-neutral speculators who take price as given would attempt to purchase infinite amounts for sale in the next period. Hence, either (1) speculators just break even ($p_t = \beta E_t(p_{t+1})$), or (2) they expect to lose money from storage. In the latter case, they carry nothing into the next period ($x_t = s_{t+1} = 0$):

$$x_t \geq 0, p_t - \beta E_t p_{t+1} \geq 0, \text{ and } x_t[p_t - \beta E_t p_{t+1}] = 0. \tag{2}$$

In addition, the current price adjusts so that the available supply matches the sum of current

and future uses:

$$D(p_t) + x_t = s_t + g + \alpha_i. \tag{3}$$

In the agricultural price support literature, α_i is interpreted as the realized supply shock to harvest g in state i; in the cap-and-trade literature, g is regarded as the size of a permit auction and $-\alpha_i$ is interpreted as the realized demand shock in state i. Denote the available supply after adjustment for the realized supply or demand shock as A_t. That is, $A_t = s_t + g + \alpha$. Hence we can re-write the market-clearing condition (equation (3)) as:

$$D(p_t) + x_t = A_t. \tag{4}$$

After the available stock is realized in period t, the available stock in the next period is still uncertain

and we denote it $\tilde{A}_{t+1}$. Given $D(\cdot), \beta \in (0, 1]$ and $E_t p_{t+1}(\tilde{A}_{t+1})$, equations (4) and (2) can be solved for p_t and x_t. Since these solutions depend on the realized available stock at t, we write them as: $p_t(A_t)$ and $x_t(A_t)$. The expected price in the previous period can then be computed:

$$E_{t-1} p_t(\tilde{A}_t) = \sum_{i=1}^{N} \pi_i p_t(A_t^i), \text{ where } A_t^i = s_t + g + \alpha_i. \tag{5}$$

Since one can deduce the expected price function from its successor, one can work backward recursively—provided one assumes that the expected price function in some distant period is given.

In reviewing the "no floor" case, we examine the sequence of expected price functions directly rather than the sequence of "value functions" generated utilizing the dynamic programming approach of Gustafson and Samuelson since our approach, unlike theirs, remains useful after government policies are introduced.

We are interested in whether a price floor inserted strictly below the initial equilibrium price in Gustafson-Samuelson's model paradoxically raises the initial price. We develop a sufficient condition for "action at a distance" regardless of whether the price floor is hard or soft. When this condition holds, we show that the hard floor raises the equilibrium price by at least as much as the soft floor does.

We now formally define each type of price floor. Suppose that in the case of "no price floor" the government auctions g units per year. We denote the "no floor" case with a superscript "N." Under either type of price floor, the auction of g units may be reduced when the price reaches the floor. When we wish to refer to the net amount sold under the auction reserve price policy (a soft floor), we append a superscript "S" to the policy function. When we wish to refer to the net amount sold (which may be negative) under the **b**uyback policy (a hard floor), we append a superscript "H" to the policy function. When we do not wish to distinguish between the two policies, we drop the superscript from the policy function $g(p)$.

Under an auction reserve price:

$$g^S(p_t)\begin{cases} = g & \text{if } p_t > f \\ \in [0,g] & \text{if } p_t = f \\ = 0 & \text{if } p_t < f \end{cases} \tag{6}$$

That is, if $p_t < f$ an auction reserve price reduces the number of units sold to zero. However, with a soft floor nothing prevents the market price from falling below f.

Instead of imposing an auction reserve price, the government may offer to buy at the per unit price f whatever amount speculators offer to sell. Since the government soaks up any excess supply offered at f per unit,

$$g^H(p_t)\begin{cases} = g & \text{if } p_t > f \\ \in [-\infty,g] & \text{if } p_t = f \end{cases} \tag{7}$$

With a hard floor, the market price can never fall below f.

Although the three policies (no floor, a soft floor at f and a hard floor at f) *seem* qualitatively different, each can be regarded as a special case of the same policy. Under that policy, the government first auctions g permits with no reserve price and then offers to buyback at price f up to the limited amount of $\bar{g}$ units. Let $R(p_t;\bar{g})$ denote the government's removal function (alternatively one might think of it as a government demand curve):

$$R(p_t;\bar{g})\begin{cases} = 0 & \text{if } p_t > f \\ \in [0,\bar{g}] & \text{if } p_t = f \\ = \bar{g} & \text{if } p_t < f \end{cases} \tag{8}$$

If there is "no floor" (N), $\bar{g} = 0$; if there is a soft floor (S), $\bar{g} = g$; and if there is a hard floor (H), $\bar{g} = \infty$.[7]. To distinguish these three policies, we write $R^j(p_t;\bar{g})$ for $j = N, S, H$.

Hence, to incorporate no price floor, a soft floor or a hard floor, one can replace equation (4) with:

$$D(p_t) + R^j(p_t,\bar{g}) + x_t = A_t. \tag{9}$$

To solve the model, we simply replace the market-clearing equation (4) in the Gustafson-Samuelson model with equation (9). Given $D(\cdot), \beta \in (0,1]$ and $E_t p^j_{t+1}(\tilde{A}_{t+1})$, equations (2), (8) and (9) can be solved for $p^j_t(A_t)$ and $x^j_t(A_t)$. But since the expected price in the previous period is

$$E_{t-1}p^j_t(\tilde{A}_t) = \sum_{i=1}^{N} \pi_i p^j_t(A^i_t) \tag{10}$$

these three equations can again be deployed to compute $p^j_{t-1}(A_{t-1})$ and $x^j_{t-1}(A_{t-1})$ and so on. Eventually, these sequences converge uniformly to the unique limit functions $x^j(A)$ and $p^j(A)$ in-

[7]We are indebted to Christophe Gouel for this insight.

dependent of the price function assumed at the outset of the process.[8] We refer to these as the infinite-horizon price and carryover functions.

3 Action at a Distance in the Infinite-Horizon Model

In this section, we develop a sufficient condition for "action at a distance" to occur. We also show that, when either type of floor would displace the equilibrium, the hard floor would raise the initial price by a weakly larger amount.

3.1 A Sufficient Condition for Action at a Distance

We focus on the stationary solution to the infinite-horizon problem with no floor (denoted by a superscript N): $p^N(A)$ and $x^N(A)$. Two properties of these limit functions will be used repeatedly: when $x^N(A) > 0$, $\frac{dx^N}{dA} \in (0,1)$ and $\frac{dp^N}{dA} < 0$.[9] Intuitively, if availability is large enough that positive carryover occurs, increasing availability by one unit will cause both carryover and consumption to increase strictly; the price must fall to stimulate increased consumption.

Price and carryover are functions of realized availability and availability evolves stochastically:

$$\tilde{A}_{t+1} = A_t - D\left(p^N(A_t)\right) + g + \tilde{\alpha}. \tag{11}$$

Since $x^N(A_t) = A - D\left(p^N(A_t)\right)$, a sequence of largest supply shocks $(\alpha_1, \ldots \alpha_1)$ will generate a deterministic path of availabilities satisfying the following first-order difference equation:

$$A_{t+1} = x(A_t) + g + \alpha_1. \tag{12}$$

Availability will converge to a unique level (denoted A^*) since the right-hand side of equation (12) is strictly positive and increases at a rate less than 1.

Define $\hat{A}$ as the unique availability level that would result in the price equalling f, the level of the hypothetical floor under consideration: $p^N(\hat{A}) = f$. Since we assume that the current price is at least as high as the floor, $A(0) \leq \hat{A}$. Logically, the fixed point A^* must be: (1) in between A_0 and $\hat{A}$; (2) smaller than A_0; or (3) larger than $\hat{A}$. In case (1), a succession of the largest harvests would not drive availability from A_0 down to A^* and hence price would never be driven down far enough to trigger government intervention if a floor of either type were imposed at f. In case (2), A_0 is so large that, despite the succession of the largest harvests, availability would fall from A_0 and the price would rise monotonically further away from the floor. In case (3), however, a succession of the largest harvests would drive availability toward the fixed point and it would inevitably exceed $\hat{A}$.

[8] See Salant (1983) for details.

[9] These properties are most easily established from the equivalent planning problem: $V(A) = \max_{x \in [0,A]} U(A - x) + \beta EV(x + g + \tilde{\alpha})$ where U is a strictly concave utility function and V is the value function, the fixed point of the contraction mapping. By a standard argument, it can be shown that V is strictly concave. When $x > 0$, $x(A)$ is defined implicitly by the first-order condition $U'(A - x) = \beta EV'(x + g + \tilde{\alpha})$. Differentiating, we conclude that $\frac{dx}{dA} = \frac{U''(A-x)}{U''(A-x)+\beta EV''(x+g+\tilde{\alpha})} \in (0,1)$. Since $p(A) = U'(A - x)$, $\frac{dp}{dA} < 0$.

Figure 1: A Sufficient Condition for a"Nonbinding" Price Floor to Raise the Current Price

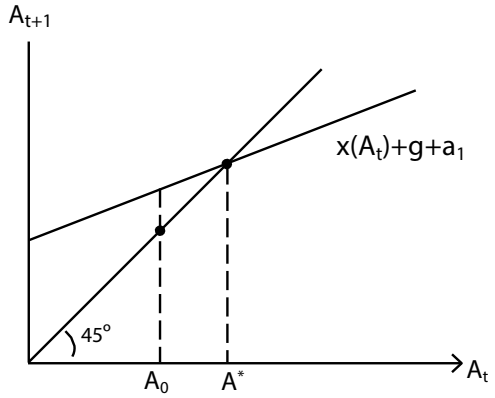

Figure 1 illustrates case 3. If threshold $\hat{A}$ is *between* A_0 and A^*, then the threshold will be crossed. Once availability exceeds $\hat{A}$, the price function $p^N(A)$ would no longer clear the market if either type of floor had been implemented since, in addition to private demand, there would be government demand. For the excess demand to be eliminated requires the price associated with availability above $\hat{A}$ to rise.

If storage occurred at the outset ($x^N(A_0) > 0$), it would continue to occur in every period until the threshold is crossed since $\frac{dx^n}{dA} > 0$. As a result, the price in the initial period must equal the discounted price expected initially to occur at every subsequent date—including those after the threshold is crossed. Consequently imposition of a floor at f would also cause the the initial price to rise, creating action at a distance.[10]

Proposition 1. *Action at a distance will occur if* $x^N(A_0) > 0$ *and* $\hat{A} \in [A_0, A^*)$.

Proposition 1 clarifies how the result of Dwight Lee (1978) generalizes to a stochastic setting. Lee considered the effects of imposing a hard ceiling in a deterministic Hotelling model above the current price (and below the choke price if one exists). A hard ceiling would change the stationary demand curve into one where the demand facing private extractors is zero at prices strictly above the ceiling since consumers would never pay more than the ceiling price for the good. If such a ceiling were imposed, it would displace the equilibrium price vector since, although the old extraction path would still maximize wealth along the old price path, it would have resulted in excess supply at prices above the ceiling. The price reduction communicates itself to the initial price since storage underground occurs from the outset.

In any stochastic model, a proposed policy may not appear to bind initially. But it will displace the equilibrium if there is any future state where the the market would no longer clear after the policy is imposed and that state occurs with positive probability. If at each stage along such a path there is storage, then the proposed policy must affect the initial price even if the policy appears "nonbinding." This stochastic generalization of Lee's result applies not only to price floors but also

[10]The condition is not necessary since action at a distance can also occur if $x^N(A_0) = 0$.

to price ceilings.

3.2 Comparing the Effects on the Current Price of Imposing Hard or Soft Floors

We now show that, in cases where there would be action at a distance, the price would be displaced at least as much if a hard floor were inserted instead of a soft floor set at the same level.

It will be helpful to denote the price in any period under policy i ($i = N, S, H$) if realized availability is A and x is carried into the next period as $P^i(A - x)$. Since the market must clear in every period, $P^N(A - x) = D^{-1}(A - x)$, $P^H(A - x) = \max\left(f, D^{-1}(A - x)\right)$ and

$$P^S(A-x) = \begin{cases} D^{-1}(A-x) & \text{if } D^{-1}(A-x) > f \\ f & \text{if } D^{-1}(A-x-g) \geq f \geq D^{-1}(A-x) \\ D^{-1}(A-x-g) & \text{if } D^{-1}(A-x-g) < f. \end{cases}$$

In a finite-horizon problem, nothing is carried out of the last period ($x_0^i = 0$). Therefore, we can easily compute $p_0^i(A) = P^i(A)$. Note that (1) $p_0^i(A)$ is bounded, (2) continuous, and (3) weakly decreasing with horizontal segments only when $p_0^i A) = f$. Moreover, (4) $p_0^N(A) \leq p_0^S(A) \leq p_0^H(A)$.

We can work backwards to determine $p_j^i(A)$ for $i = N, S, H$ and $j = 0, 1, \ldots$ and can verify that each of these functions share these four properties.

For a given policy i and final price function $p_0^i(A)$, suppose we enter one period back with some realized availability A, carry x_1^i into the last period and expect that each unit carried into the final period would earn $\beta \sum_k \pi_k p_0^i(x_1^i + g + \alpha_k) = \beta E p_0^i(x_1^i + g + \tilde{\alpha})$. Then x_1^i uniquely solves the following complementary slackness condition:

$$x_1^i \geq 0, P^i(A - x_1^i) - \beta E p_0^i(x_1^i + g + \tilde{\alpha}) \geq 0, \text{ with complementary slackness.}$$

Denote the solution as $x_1^i(A)$. Market clearing requires that $p_1^i(A) = P^i\left(A - x_1^i(A)\right)$. Note again that $p_1^i(A)$ inherits the four properties.

This process of deducing $p_1^i(A)$ from $p_0^i(A)$ can be applied to any price function $(p_j^i(A))$ with these four properties to obtain $p_{j+1}^i(A)$. We summarize the process by the operator T_i for $i = N, S, H$. That is, $p_{j+1}^i(A) = T_i p_j^i(A)$. Since the operator satisfies the Blackwell conditions, there exists a unique fixed point, denoted $p^i(A)$ (with no subscript). That is, $p^i(A) = T_i p^i(A)$. If the operator is applied repeatedly to a starting function with the four properties, its image will converge uniformly to the fixed point, $p^i(A)$. The fixed point inherits the four properties.

If the current realized availability is A and x is carried into the next period, then the discounted expected price next period will be $\beta \sum_k \pi_k p^i(x + g + \alpha_k)$. Hence, the expected discounted price function is weakly decreasing with horizontal segments at βf.

Suppose no floor is in place ($i = N$) and the period begins with stock s carried in. Suppose shock $\alpha_k << \alpha_1$ occurs. Denote the realized availability following α_k as $A_k = s + g + \alpha_k$. Assume $p^N(A_k)$

Figure 2: Simultaneous Determination of Carryover ($x^i(A_k)$) and Price ($p^i(A_k)$) When Unanticipated Floor Is Inserted Just Below the Spot Price

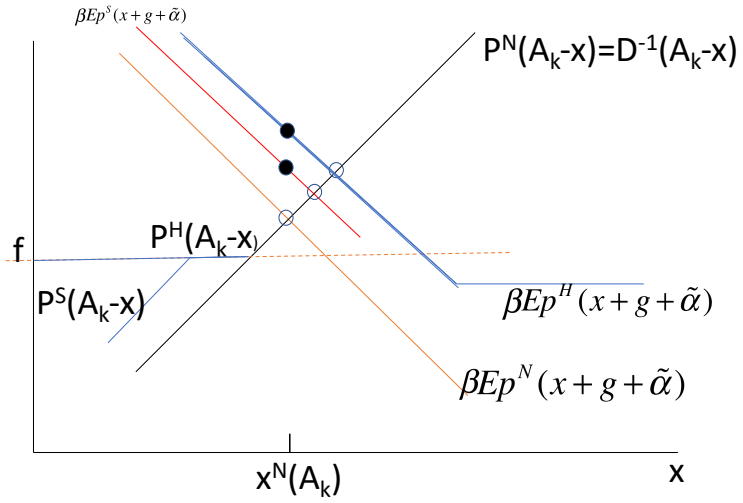

is marginally higher than f. We depict the situation in Figure 2. The downward-sloping curve is $\beta Ep^N(x^N+g+\tilde{\alpha})$. The upward-sloping curve is $P^N(A_k-x^N)=D^{-1}(A_k-x^N)$. By assumption, the vertical component of the intersection point just exceeds f. The horizontal component determines $x^N(A_k)$. If a soft or hard floor is imposed, it will be "nonbinding" since the price with either type of floor will be at least as high as the price with no floor when evaluated at the same realized availability (A_k) and hence will strictly exceed f.

Suppose, however, that for the *largest* supply shock, α_1, $p^N(x^N+g+\alpha_1)<f$. This can easily happen since shock $\alpha_k << \alpha_1$ and $p^N(x^N+g+\alpha_k)$ is close to f. In that case, $p^H(x^N+g+\alpha_1)>p^N(x^N+g+\alpha_1)$ since p^H is never below f $(p^H \geq f)$ Since in every other state k the price with a hard floor is weakly higher, it follows that $Ep^H(x^N+g+\tilde{\alpha})>Ep^N(x^N+g+\tilde{\alpha})$. Indeed, since the price function p^N is weakly decreasing, $p^N(y+g+\alpha_1)\leq p^N(x^N+g+\alpha_1)<f$ for $y\geq x^N$. Therefore, $Ep^H(y+g+\tilde{\alpha})>Ep^N(y+g+\tilde{\alpha})$ for *any* $y\geq x^N$.

Suppose that for the largest shock, α_1, $p^S(x^N+g+\alpha_1)<f$. This is not inevitable but is certainly possible. In that case, $p^H(x^N+g+\alpha_1)>p^S(x^N+g+\alpha_1)>p^N(x^N+g+\alpha_1)$. It follows that $Ep^H(x^N+g+\tilde{\alpha})>Ep^S(x^N+g+\tilde{\alpha})>Ep^N(x^N+g+\tilde{\alpha})$. Indeed, since the price function $p^S(x^N+g+\alpha_1)$ is weakly decreasing, $p^S(y+g+\alpha_1)\leq p^S(x^N+g+\alpha_1)<f$ for $y\geq x^N$. Therefore, for *any* $y\geq x_N$, $Ep^H(y+g+\tilde{\alpha})>Ep^S(y+g+\tilde{\alpha})>Ep^N(y+g+\tilde{\alpha})$.

Figure 2 illustrates the case where (1) the price with no floor is marginally higher than f and (2) $p^S(x^N+g+\alpha_1)<f$ and so the three expected price functions can be strictly ranked for carryover x^N : $Ep^H(x^N+g+\tilde{\alpha})>Ep^S(x^N+g+\tilde{\alpha})>Ep^N(x^N+g+\tilde{\alpha})$.[11] Consequently, $x^H(A_k)>x^S(A_k)>x^N(A_k)$ and $p^H(A_k)>p^S(A_k)>p^N(A_k)$.[12] The case where $p^H(A_k)=p^S(A_k)$

[11]The case where $p^H(A_k)=p^S(A_k)$ and $x^H(A_k)=x^S(A_k)$ is not depicted.

[12]An alternate way of deriving these results follows from footnotes 7 and 9. The carryover function for the general

and $x^H(A_k) = x^S(A_k)$ is not depicted. Intuitively, since for a given carryover of x^N or larger the discounted expected price is strictly higher with a hard floor than with no floor, more is carried out of the current period if the floor is hard, driving the current price above what would it would be in the absence of any floor. A similar intuition explains why $x^H(A_k) > x^S(A_k) > x^N(A_k)$ if $Ep^H(x^N + g + \tilde{\alpha}) > Ep^S(x^N + g + \tilde{\alpha}) > Ep^N(x^N + g + \tilde{\alpha})$.

We conclude therefore:

Proposition 2. *Introducing a price floor under the current spot price (but close to it) will cause the spot price to jump up and that the jump may be higher if the floor is hard rather than soft.*

.

To illustrate our key points, we carried out a set of simulations of the three policy regimes. We used Gouel's (2017) very useful tool for specifying, solving and simulating rational expectations equilibria with the Matlab CompEcon tools developed by Miranda and Fackler (2002). We used a linear demand function similar to the one used in our experimental investigations. The annual new supply, g, is constant. Each period α_i is either high or low, which occur with equal probability. The policies of no floor, soft floor and hard floor are implemented as maximum buyback amounts ($\bar{g}$) of 0, g, or ∞ at a given floor price. We calculate the rational expectations equilibrium carryover and price functions: $x(A)$ and $p(A)$.[13]

Figure 3 shows how expected price varies with the stock (x) carried into the period under the three policy regimes with a \$105 floor price. The expected price functions for the soft and hard price floor policies have the predicted relationship.

Figure 4 shows the current equilibrium price as a function of availability, again using a price floor of \$105. Prices under both the soft and hard floor regimes lie uniformly above the no floor policy. At any availability below 36, the floor is not binding in any of the three policy regimes, and yet, if availability is slightly smaller than 36 and hence the price is marginally above the floor price of \$105, the price of the soft floor is higher than the no floor case, and the hard floor price is higher still. Thus, if an unanticipated floor of \$105 were imposed when the price was slightly above that level, the price would jump up with a soft floor and would jump higher with a hard floor—an illustration of *action at a distance.*

To show the effect of the price floor level on prices, we ran 1,000 simulations of the model at different price floors and calculated the average price for each policy regime in each of the 200 periods the simulation runs. Figure 5 give the results for four different price floor levels: \$55, \$80, \$100 and \$110. At the \$55 price floor, all three policy regimes are indistinguishable, since the floor is essentially never binding. The \$80 price floor induces a separation between the no floor policy and

policy of auctioning g units and then offering to buy back up to $\bar{g}$ units at price f can be deduced from the following dynamic programming problem: $\hat{V}(A) = max_{x\in[0,A]} \int_{q=0}^{q=A-x} \hat{U}'(q)dq + \beta E\hat{V}(x+g+\alpha)$ where $\hat{U}(\cdot)$ is a *pseudo* utility function defined as follows: $\hat{U}'(q) = D^{-1}(q)$ for $q < q_1$ or $q > q_1 + \bar{g}$ and $\hat{U}'(q) = f$ for $q \in [q_1.q_1 + \bar{g}]$. where $q_1 = D^{-1}(f)$. After computing $x(A;\bar{g})$ from this dynamic programming problem, $p(A;\bar{g}) = P^i\left(A - x(A;\bar{g})\right)$ can be computed. Finally, $\bar{g}$ can be set at $0, g,$ or ∞ for the three policies of interest.

[13]For our simulations, we use a demand intercept of 105, slope of 0.75, g of 40, discount factor of 0.95 and shock size of 50.

Figure 3: Expected Current Price Functions Under No Policy, a Soft Floor and a Hard Floor in the Linear Demand Example

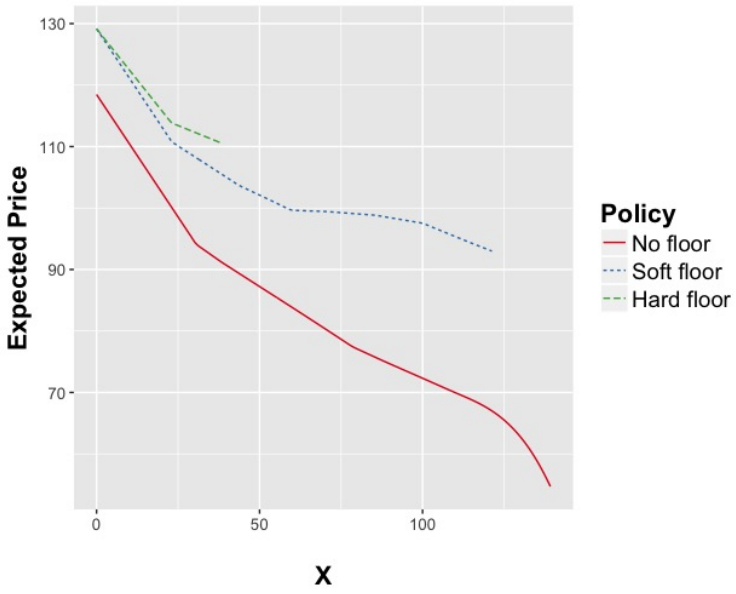

Figure 4: Current Price as a Function of Availability

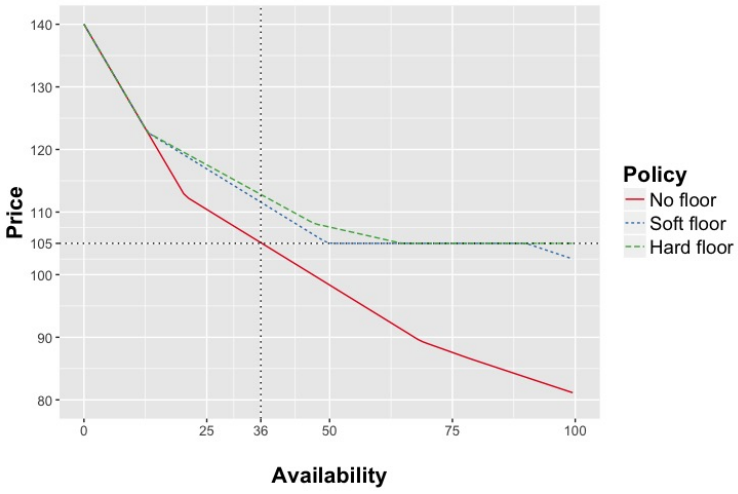

Figure 5: Average Prices for Different Price Floors and Policy Regimes

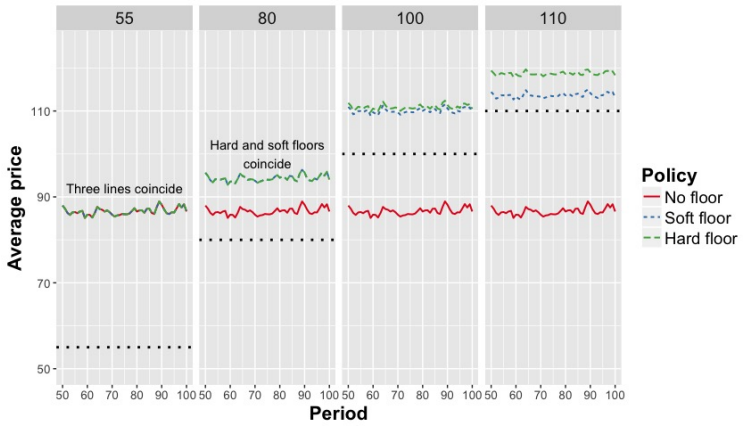

the other two policies, which are indistinguishable from each other. At progressively higher price floors, the average price under the hard floor moves higher than for the soft floor, and this difference increases as the floor rises.

3.3 A Comparison to Krugman's Target Zone Model

In his classic article on target zones, Krugman (1991) discusses the behavior of a foreign exchange rate supported by the same hard floor policy as us. Moreover, a unit of foreign exchange is storable, just like a unit of grain or an emissions permit. But here the similarity ends. Our assets disappear at a price sensitive rate: the lower the price, the more grain is consumed and the more permits are surrendered at true up. That is why an auction is needed in our model to replenish the assets while no auction is needed in Krugman's model. Foreign exchange, on the other hand, is held for its "convenience yield," like a painting.[14] These economic differences, rather than the fact that Krugman (1991) conducts his published analysis in continuous time,[15] account for the two striking differences between his conclusions and ours.

First, unlike the "fundamental" (velocity) in Krugman's target zone model which shifts the demand for services of the asset, our fundamental (availability) does not execute an additive random walk with zero drift. In ours, $\tilde{A}_{t+1} = A_t - D\left(p^i(A_t)\right) + g + \tilde{\alpha}$. We conclude, therefore:

Proposition 3. *Expected drift, conditional on current availability, is:*

$$E(A_{t+1}|A_t) = g - D\left(p^i(A_t)\right). \tag{13}$$

Therefore, expected drift is price-sensitive. Even if it were zero for some availability level, it would be negative for anything larger and positive for anything smaller. In the absence of a floor,

[14] Indeed it would be an interesting exercise to derive results for the painting price using a procedure like the one we have followed.

[15] Krugman (1987) presents a discrete-time analysis.

Figure 6: The Hard Floor Does Not Induce Smooth Pasting in Our Model

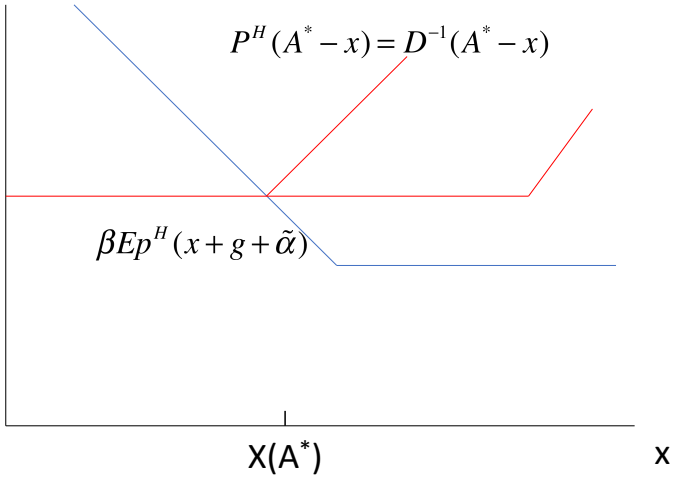

Krugman's foreign exchange rate will eventually hit any proposed floor. Our price, however, may never reach a proposed floor if it is set too low. For that reason, we had to work out a condition sufficient for the floor to be hit with positive probability.

Second, in Krugman's model, the exchange rate is insensitive to changes in the fundamental when the exchange rate reaches the floor—his "smooth pasting" result. In our model, the left derivative of $p^H(A)$ is not zero as in his model. Let A^* denote the smallest availability generating $p^H(A) = f$. The left derivative, $p^H(A^*)$is strictly negative. This can be seen in our Matlab results of Figure 4 and can be proved analytically as follows. Consider an availability level just large enough to drive the price down to the hard floor. We depict this situation in Figure 6.

Proposition 4. *There is no smooth pasting result in our model*

Proof. We want to calculate the left derivative, $\frac{dp^H}{dA}$ evaluated at A^*. Since $p^H(A) = D^{-1}(A - x(A))$, the left derivative is strictly negative if the left derivative $\frac{dx^H}{dA} < 1$ when evaluated at A^*. Now the carryover at this point is implicitly defined as the solution to $D^{-1}(A - x) = \beta E p^H(x + g + \tilde{\alpha})$. Differentiating, we conclude that the left derivative is:

$$\frac{dx^H}{dA} = \frac{D^{-1\prime}(A - x) + E p^{H\prime}(x + g + \tilde{\alpha})}{D^{-1\prime}(A - x)} < 1. \tag{14}$$

Hence, in our model there is no smooth pasting result. □

4 Action at a Distance in the Two-Period Model

In this section, we develop an undiscounted, two-period stochastic storage model to isolate the phenomena we will be investigating in the laboratory: (1) action at a distance and (2) the stronger

effect of the hard floor. We assume here that the shocks affect demand rather than supply since that is the formulation we used in the experiment. The section concludes with an example. After presenting that example, we show how it can be reinterpreted if the demand curve is stationary and the shocks are to supply.

Assume the demand curve in each period is $Q = a - p/m$; so the inverse demand is $p = m(a - Q)$. Supply in the first period is exogenous: A_0. This is either consumed in the first period or carried into the second period (denoted x).

$$\overbrace{a - p/m}^{\text{consumption}} + \overbrace{x}^{\text{carryover}} = \overbrace{A_0}^{\text{supply}}. \tag{15}$$

Hence, the first-period price as a function of the carryover is:

$$D^{-1}(A_0 - x) = m(a - (A_0 - x)). \tag{16}$$

Substituting into the demand function,we conclude that the first-period consumption is $Q = A_0 - x$.

In the second period, a can take on two values: $a_h > a_l$. Supply in the second period consists of what is carried over from the first period (denoted x) plus what the government distributes in the second period (denoted g). If there is no price floor (denoted superscript N), the government purchases nothing. Demand in the second period consists only of consumption demand since nothing is carried out of the last in the absence of a bequest motive.

Supply in the second period is $x + g$. Demand in state k ($k \in h, l$) is $a_k - p_k/m$.

$$\overbrace{a_k - p_k/m}^{\text{consumption}} = \overbrace{x}^{\text{carryover}} + \overbrace{g}^{government}, k = h, l. \tag{17}$$

Hence, the second-period price as a function of the carryover in the absence of a price floor is state-dependent:

$$p_k^N(x) = m(a_k - x - g). \tag{18}$$

Substituting into the demand function, we conclude that the second-period consumption is $Q_k = x + g$. That is, the price in each state adjusts so that everything available in the second period is consumed.

Hence, the price expected in period 1 to prevail in period 2 is a decreasing function of the carryover. If there is no floor:

$$Ep^N(x) = m\left(\pi_h[a_h - g - x] + \pi_l[a_l - g - x]\right) = m\left(E(a) - x - g\right). \tag{19}$$

Assume that $D^{-1}(A_0) < Ep^N(0)$ (which requires $a - A_0 < E(a) - g$) and that carryovers are costless (zero interest rate and zero storage cost). Then permits will be carried into the second period until $D^{-1}(A_0 - x) = Ep^N(x)$ for some $x > 0$. This equation defines the equilibrium carryover (x^N) when there is no floor.

Suppose we have a price floor f. Let $\hat{x}$ be the threshold carryover that would result in price in the low-demand state just driven to the floor: $p_l(\hat{x}) = m(a_l - x - g) = f$.[16]

$$\hat{x} = a_l - g - f/m. \tag{20}$$

If the equilibrium carryover were *strictly smaller* than $\hat{x}$, the state-dependent price in the second period would exceed the floor even in the low-demand state, and the expected price in that period with either type of floor would coincide with the expected price in the absence of a floor. Hence, the equilibrium price in the first period would not be affected by the imposition of the floor. This case is not very interesting!

However, if the equilibrium carryover with no floor exceeds $\hat{x}$, then the equilibrium price in the low-demand state will be affected when a floor is imposed. If the price floor is hard (denoted superscript H), then the price in the low-demand state will be f. Consequently, for $x > \hat{x}$,

$$Ep^H(x) = m\pi_h(a_h - g - x) + \pi_l f > Ep^N(x). \tag{21}$$

Note that, with a hard floor, a unit increase in carryovers for $x > \hat{x}$ drives the expected price down by only $m\pi_h < m$ since it has no effect on the price in the low-demand state.

If the price floor is soft (denoted superscript S), then the price in the low-demand state depends on whether, in equilibrium, $x^S > \hat{x} + g$ or $x^S \in (\hat{x}, \hat{x} + g)$.

In the former case, the price in the low-demand state will be uniformly strictly smaller than f and therefore strictly smaller than with a hard floor. As a result, the expected price with a hard floor will be uniformly above the expected price with a soft floor for $x^S > \hat{x} + g$.

Figure 7 depicts the former situation. The horizontal axis has length A_0. Carryovers are measured along the horizontal axis from the origin rightward up to A_0. The first-period price function $D^{-1}(A_0 - x)$, depicted in green and defined in equation (16)), increases at rate m as more is carried out of the first period. $Ep^N(x)$, depicted in red and defined in equation (19), decreases at rate m as more is carried into the second period. The two functions intersect at (x^N, p^N), where $x^N > \hat{x} + g$. The price floor (not depicted) would be marginally below p^N, the vertical component of this point of intersection.

The expected price in period 2 with a hard floor is piecewise linear. It has *two* segments: a red segment and a blue segment. Its red segment coincides with $Ep^N(x)$ for $x \leq \hat{x}$ but then kinks to a flatter blue line for any $x > \hat{x}$. It intersects the first-period price function northeast of the previous intersection point. That is, in equilibrium speculators carry out of period 1 more permits when protected against a downside price shock by the hard floor. Consequently the first period price is higher.

The expected price in period 2 with a soft floor is also piecewise linear, but with *three* segments instead of two. It has a red segment, a blue segment and a black segment. If the carryover is smaller than $\hat{x}$ then this expected price function coincides with the other two expected price functions. If

[16] Note that in the high-demand state, that carryover would result in a price above f.

Figure 7: Differential Effects of Hard and Soft Floors

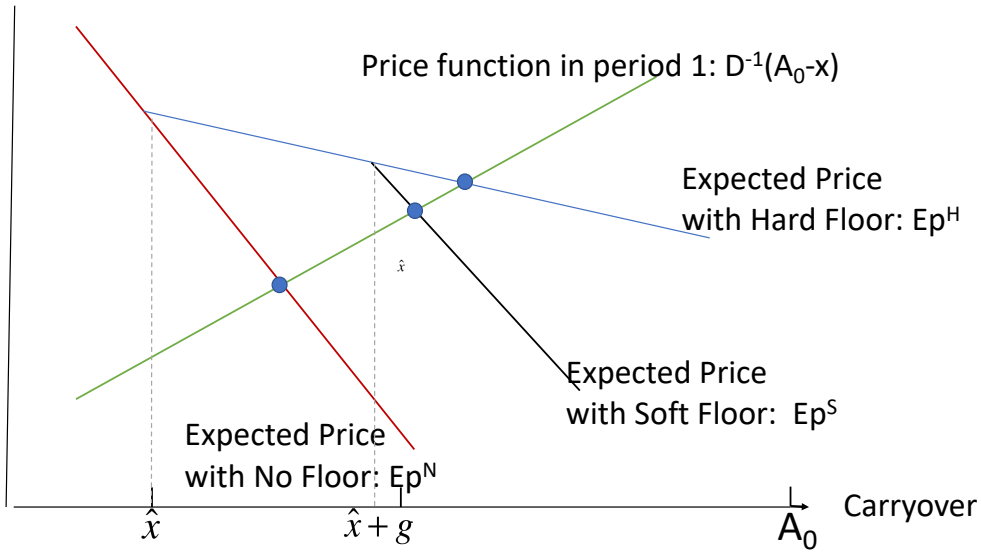

$x \in (\hat{x}, \hat{x}+g)$ then the expected price function with the soft floor coincides with the expected price function with the hard floor, decreasing at rate $m\pi_h$. If, however, the equilibrium carryover exceeds $\hat{x}+g$, then the additional carryover will not be absorbed by the government since the government is constrained not to buyback more than g (or, equivalently, constrained to auction zero if no bid is as high as the auction reserve price). Hence the expected price function declines linearly at rate m (so the black segment is parallel to the red segment).

As the diagram illustrates, the price floor affects the price expected to prevail in the next period and hence the price and carryover in the current period even though $p^N > f$: so-called "action at a distance." Since the government in the hard floor case is unconstrained in how much it can buy back to maintain the floor, it will raise the price by more than a soft floor provided $x^N > \hat{x}+g$.

If A_0 were smaller, the green line in Figure 7 might intersect the blue line with a horizontal component in the interval $(\hat{x}, \hat{x}+g)$. In that case, the soft and hard floors would have identical effects on price in the first period, on carryovers into the second period and on the state-dependent prices in the second period.

In our experimental design, however, we wanted the two floor policies to have distinguishable effects. Therefore, we set the floor so that:

1. it is marginally below the equilibrium price in the first period
2. it does not bind in the high-demand state
3. it induces a carryover larger than $\hat{x}+g$ so that the hard and soft floors are distinguishable.

We conclude this section with a parametrized two-period example where the first period price jumps when a "non-binding" floor is imposed and, given the height of the floor, the jump is larger with a hard price floor than with a soft one.

The example is summarized in Table 1. The example is described in terms of demand shocks, the formulation used in our laboratory experiment.

Demand in each period is linear with slope $-.5$. In the second period demand shifts left or right with equal probability without changing slope. Agents are initially endowed with 40 units of the good and can augment anything they carry into the second period with acquisitions from the auction of 8 units.

Table 1 indicates our predictions.[17] In the absence of a price floor, the equilibrium carryover is 16 units and the first-year price is 152. As the last two rows of the table indicate, neither type of floor will affect the carryover or prices if the floor is set extremely low (68 in the table). However action at a distance occurs for higher floors. With a floor of 128, the effects of the two floors differ: the price jumps to 156 with a soft floor but to 164 with a hard floor.

5 A Laboratory Experiment to Test the Theory

The two period model that we have just presented gives rise to sharp predictions that we can test in the laboratory. In this section, we describe our methodology for performing these tests.

As of this writing, we are conducting experimental sessions at the AREC Experimental Economics Laboratory at the University of Maryland. We are using the z-Tree experimental software (Fischbacher, 2007). We recruit University of Maryland students to participate in sessions that last approximately an hour and a half. We recruit subjects on a first-come-first-serve basis from a large pool of potential participants representing different majors and different grade levels. The experimenter reads the instructions aloud at the beginning of each session. The instructions (see Appendix) are followed by a short quiz that tests and reinforces student understanding of the experimental setup.

We follow a between-subjects design where subjects will only participate in one treatment. In total, we plan to run 5 treatments. Each treatment will consist of 6 sessions with 10 subjects participating. Therefore, in total we will have 300 subjects.

Each session will consist of two parts: baseline rounds and experimental rounds. Both parts will have 5 rounds, 10 rounds in total. In each round, subjects participate in a 2-period market game where each of them has the role of a trader in a market for a generic commodity (grain). They will be asked to decide how much grain to buy, sell, or store.

[17] The reader can easily verify that the prediction in each row of the table is an equilibrium: (1) since the carryover is strictly positive, the price in year 1 equals the expected price in year 2; (2) the price in year 1 equates the demand in that year to the initial allocation minus the carryover; and (3) the price in year two equates the private demand in that year and state plus the government buybacks at f (constrained not to exceed $\bar{g}$) to the sum of the carryover and the auction.

Table 1: Predictions in a Two-Period Example

	Floor	Carry Over x	Price in Year 1 P_0	Year 2 Price if demand is H	Year 2 Price if demand is L
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

- Given the parameters of our two-year model ($A_0 = 40, g =$ 8, D=$a_j - .5p$ f or a_0=100, a_L=70, a_H=130), our hypotheses are:
  H_1: At $P_{floor} = 68$, no effect on equilibrium price or carryforward
  H_2: At $P_{floor} = 128$, $P_N < P_S < P_H$

Upon completion of the experiment, one round from Part 1 and one round from Part 2 will be randomly chosen for payment. Subjects 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$) = 1.00 US Dollar ($). Subjects earnings (plus the $7 show-up reward) will be paid to them in private at the end of the experiment. We expect subjects earnings will be in the range $20-$30.

In the beginning of the experiment, subjects will be told that there will be 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 will be given a quiz to check for understanding.

Part 1 of the experiment will not have any price floor and will be identical among all our treatments. Part 1 will serve two main purposes. First, it will allow subjects to learn about the procedures of the experiment. Second, and more importantly, it will let 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.

Part 2 of our experiment will differ across treatments. Our control treatment (BASELINE) will repeat the same 2-year market game as in part 1 for another 5 rounds. Soft floor treatments will impose a reserve price in the auction, while hard floor treatments will have a buyback guarantee in the spot market. The baseline treatment has the added advantage of controlling for any restart effects subsequent to the break at the conclusion of Part 1.

At the end of Part 1, participants will be told that there will be another 5 rounds of the game and will be told about any changes to the rules from Part 1. So, Part 2 starts with a brief announcement about the instructions for Part 2. In the BASELINE treatment participants are told that the instructions are the same as for Part 1. In the other treatments the subjects are told either about the auction reserve price or the guaranteed buyback at a set price. Recall that in both parts, subjects will face the same 2-year market game.

In every round, at the beginning of the first year, subjects receive 4 units of grain 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 their grain they own to (pre-programmed) computerized *consumers*. In the first year, consumer demand for grain is: $Q = 100 - 0.5P$.

As sellers, participants will be 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 year is automatically carried into the second year. At the end of the second year, any unsold grain will have no value.

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 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.

At the start of the second year, 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 8 units of grain are available in a sealed-bid, uniform price auction. Subjects will be able to bid (anonymously and simultaneously) for up to two units, subject to their bids not exceeding their current cash balance. The bids will be ranked from high to low. The asset being auctioned will be sold to the highest bidders at a uniform price that is the value of the highest rejected bid, so all winning bidders will pay the same price. If there are 8 or fewer bids, then all bids will be successful and the grain will be free. Any unsold grain will no longer be available to the market. After the grain auction, subjects will again participate in the spot market as sellers. In the second year, demand can either be high (Q=130-0.5P) or low (Q=70-0.5P) with equal probabilities. Subjects are told how the spot price is determined. In order to make it easy for subjects, instead of giving formulas, we will provide them with a table of the valuations of the computerized buyers (see Table A1, in the appendix).

In our soft floor treatments (SOFT), subjects are told at the beginning of Part 2 that there is one change compared to Part 1; there will be a reserve price in the auction. Subjects will not be able to submit any bids that are lower than the reserve price. Similarly, in our hard floor treatments (HARD), at the beginning of Part 2 subjects will be told that there is one change compared to Part 1; there will be a special computerized buyer willing to buy unlimited amounts at a certain price (our hard floor price) in year 2. We will implement the treatments with two different price floors. In SOFT-HIGH and HARD-HIGH treatments, the price floor will be 128. The price floor in the SOFT-LOW and HARD-LOW equals 68.

Theoretical predictions corresponding to our different treatments are presented in Table 1. Our experiments will test whether the comparative static predictions based on our model hold in the data. Note that, in order to test our predictions, we will rely on data from part 2 of each treatment.

6 Conclusions

In this paper, we explicitly model firm responses to non-binding price floors in the emission market context. We have shown that seemingly non-binding hard and soft price floors can be expected to elevate the current spot price. This amplifies results left implicit in existing literature on the effects of non-binding price controls and confirms what Holt and Shobe (2016) observed in their laboratory experiment with soft price floors and what the *Financial Times* (2013) observed in Chinese rice markets with hard floors. We also show that the two kinds of floors differ in their consequences. This is particularly important because price floors in cap-and-trade programs have historically always been soft floors but virtually every non-experimental analysis of them has treated them as hard floors. Our theoretical results are borne out in simulations of the infinite-horizon rational expectations equilibrium. In addition to the infinite horizon model, we developed a two-period model We are now in the process of carrying out an experimental investigation to test the model predictions.

Appendix: Instructions to Subjects in Laboratory Experiment

Instructions:

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** 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** 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

(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. 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.

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

(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	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 A.2: How the Price is Calculated in the Grain Market

(1)	(2)	(3)
	Year 2 Valuations of Pre-Programmed Consumers if Demand is HIGH	**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 A.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	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

References

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

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

Eeckhouldt, Louis, and Pierre Hansen. 1980. "Minimum and Maximum Prices, Uncertainty, and the Theory of the Competitive Firm." *American Economic Review*, 70(5): 1064–1068.

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.

Gardner, Bruce. 1979. *Optimal Stockpiling of Grain.* D.C. Heath and Company.

Gouel, Christophe. 2017. "RECS toolbox (Version 0.7)." *Retrieved from www.recs-solver.org.*

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.

Johnson, D. Gale. 1954. "Competition in Agriculture: Fact or Fiction." *The American Economics Review*, 44(2): 107–115.

Krugman, Paul. 1987. "Trigger Straegies and Price Dynamics in Equity and Foreign Exchange Markets." NBER Working Paper No. 2459.

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.

McKibbin, W. J., and P. J. Wilcoxen. 2002. "The Role of Economics in Climate Change Policy." *Journal of Economic Perspectives*, 16(2): 23.

Miranda, Mario J., and Paul L. Fackler. 2002. *Applied Computational Economics and Finance.* MIT Press.

Miranda, Mario J., and Peter Helmberger. 1988. "The Effects of Commodity Price Stabilization Programs." *American Economic Review*, 78(1): 46–58.

Montgomery, David W. 1972. "Markets in licenses and efficient pollution control programs." *Journal of Economic Theory*, 5: 395–418.

Murray, Brian, Richard G. Newell, and William A. Pizer. 2009. "Balancing Cost and Emissions Certainty: An Allowance Reserve for Cap-and-Trade." *Review of Environmental Economics and Policy*, 3(1): 84–103.

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

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): 703718.

Pizer, William A. 2002. "Combining price and quantity controls to mitigate global climate change." *Journal of Public Economics*, 85: 409–432.

Roberts, Marc J., and Michael Spence. 1976. "Effluent charges and licenses under uncertainty." *Journal of Public Economics*, 5(3-4): 193–208.

Salant, Stephen J. 1980. "The Vulnerability of Price Stabilization Schemes to Speculative Attack." Federal Trade Commission, June.

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

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

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.

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.

Terazona, Emiko. 2013. "Beijing's Rice Subsidies Buoy Imports." *Financial Times*, , (August 19).

Weitzman, Martin L. 1974. "Prices Vs. Quantities." *The Review of Economic Studies*, 41(4): 477–491.

Williams, Jeffrey C., and Brian D. Wright. 1991. *Storage and Commodity Markets.* Cambridge Press.

Working, Holbrook. 1953. "Price Supports and the Effectiveness of Hedging." *American Economic Review*, 35(5): 811–818.

Wright, Brian D., and Jeffrey C. Williams. 1988. "The Incidence of Market-Stabilixing Price Support Schemes." *The Economic Journal*, 98(393): 1183–1198.