

BEHAVIORAL AND EXPERIMENTAL ECONOMICS RESEARCH LAB

Price distortions and hoarding: An experiment

Edgar Castro • César Martinelli • Nicholas Brown

September 2026



Price distortions and hoarding: An experiment^{*}

Edgar Castro[†]

César Martinelli[‡]

Nicholas Brown[§]

September 28, 2026

Abstract

We study how current and anticipated price distortions affect storage decisions in a two-period model. Both anticipated rationing and anticipated price liberalization create incentives to hoard, though for different reasons. We test these predictions in the laboratory using three treatments: a free market with rising costs, an anticipated price cap, and liberalization following a price cap. The model successfully predicts the broad patterns of shortages and prices across treatments. Under liberalization, shortages become increasingly severe as participants gain experience with the regulated market, and prices more than double once the cap is lifted. When a price cap is anticipated, however, participants store less than predicted and bid too little for storage, as if they underestimate the severity of future rationing. Both price-control regimes entail substantial welfare losses with respect to the free market, with liberalization producing the lowest surplus as rent-seeking rather than defensive storage becomes the dominant motive for hoarding.

Keywords: price controls, price distortions, hoarding, rationing, laboratory experiments

JEL Classification: D41, C92

^{*}An earlier version of this paper was Chapter 1 of Castro's doctoral dissertation at George Mason University. We would like to thank Mikhail Freer, Shan Gui, Ruolong Xiao, and participants at several audiences for their comments.

[†]Tecnológico de Monterrey, ecastrom@tec.mx (Corresponding author).

[‡]George Mason University, cmarti33@gmu.edu.

[§]George Mason University, nbrown38@gmu.edu.

1 Introduction

Where there are price controls, there is hoarding. In Venezuela, years of price controls and shortages led families to stockpile goods at home, even as politicians declared war on hoarders.¹ In Argentina, government-imposed price controls aimed at curbing inflation produced a similar response (National Public Radio, 2023). But hoarding is not confined to inflation-plagued economies or to the defunct Soviet Union. As Guenette (2020) documents, price controls are more prevalent than commonly believed, affecting both developed and emerging economies. With the onset of the COVID-19 pandemic, many countries experienced shortages and hoarding, and in the United States officials invoked anti-price-gouging laws against resellers (Rydzewski, 2020; Department of Justice, 2022). The pandemic also revived interest in price controls as a policy tool. In 2024, Vice President Harris proposed a federal ban on price gouging for groceries, while the Trump administration subsequently signed an executive order addressing anti-competitive behavior in the food supply chain and imposed broad tariffs that directly distort import prices.²

Despite this renewed policy attention, experimental evidence on the relationship between price distortions and hoarding remains scarce. Hansman et al. (2020) provide suggestive evidence that sticky retail prices contributed to hoarding during the 2008 rice crisis, as consumers who anticipated price increases stockpiled before prices adjusted. Observational data, however, cannot easily disentangle the mechanisms through which price distortions generate hoarding or isolate storage from other behavioral responses such as panic buying. Laboratory experiments offer a controlled environment in which to identify these mechanisms, yet only a handful of experiments have studied markets with price controls (Smith and Williams, 1981; Isaac and Plott, 1981; Coursey and Smith, 1983; Finley et al., 2019), and none have examined how price distortions interact with storage to generate hoarding.

A classic explanation for hoarding under price controls is provided by Weitzman (1991). In his account of persistent scarcity in the Soviet Union, controlled prices raise the effort

¹See, e.g., Wall Street Journal (2016). Figure 16 in the appendix shows a Caracas supermarket in 2014 with a sign reading “no more than 6 cans of tuna per person, thank you.”

²See Wilkie (2024) and The White House (2025).

required to obtain goods, giving rational consumers an incentive to stockpile. Our mechanism is different but related: rationing itself creates an incentive to store when consumers anticipate that they may be unable to purchase in the future. Rationing intended to manage scarcity can therefore induce consumers to accumulate goods in anticipation of being unable to obtain them later.³

This paper studies theoretically and experimentally how current and anticipated price distortions affect storage decisions. We model hoarding as a rational response to rationing risk. When a price cap holds the price below its market-clearing level, demand exceeds supply and buyers face a risk of being unable to purchase. This creates incentives to store goods, with consequences for equilibrium prices, quantities, and welfare. Importantly, storage can arise for two distinct reasons. Defensive storage allows buyers to insure against expected future rationing. Rent-seeking storage allows them to obtain goods at a regulated price in anticipation of selling or consuming them after prices are liberalized. Although both mechanisms generate hoarding, they have different implications for allocative efficiency.

We study three environments. In all three, buyers anticipate rising production costs. In the baseline treatment, Free Market (FM), prices are unrestricted in both periods, and intertemporal arbitrage motivates storage. In the Anticipated Price Cap (APC) treatment, the first period is a free market and a price cap is imposed in the second, creating an incentive for defensive storage in response to anticipated rationing. In the Liberalization (LIB) treatment, a price cap binds in the first period and is removed in the second, creating first-period shortages and incentives for rent-seeking storage.⁴ The model yields predictions for storage, prices, shortages, and welfare across these three environments, which we test in the laboratory.

³A vivid example occurred in Bogota during the 1997–98 El Niño drought. When authorities implemented water rationing to protect reservoir levels, rationing announcements prompted Bogotans to fill improvised water tanks at home, paradoxically increasing consumption. The crisis was ultimately resolved not by rationing but by raising water utility prices (Vallejo, 2011; Ivanova and Sarmiento, 2013). In Colombia, water supply is a regulated monopoly; the Water Regulatory Commission (CRA) modified its pricing methodology through Resolution 15 of 1997 to allow price increases in larger cities. Over the following 12 years, average water consumption fell by 17%.

⁴A historical example occurred in Argentina in 1990, when President Carlos Menem announced a 25% increase in gasoline prices associated with liberalization. The announcement triggered long queues at service stations and gasoline shortages throughout the country (Archivo DiFilm, 1991).

The model predicts the broad patterns observed in the experiment, but we also find systematic departures from equilibrium. Most notably, participants store too little when storage must anticipate future rationing. In the Anticipated Price Cap treatment, participants store 2.3 units per market rather than the predicted 3, and the first-period price averages 73 rather than the predicted range of [88,90]. Participants also bid too little for storage, behaving as if they underestimate the tightness of the future rationed market and expect to be able to purchase there. By contrast, in the Liberalization treatment, participants increasingly compete for goods at the controlled price. The shortage rate in the capped first period rises from 12% in round 1 to 34% in round 10, while the second-period price more than doubles to 109 once the cap is removed. Periods without a price cap have near-zero shortage rates in all three treatments.

These differences translate into substantial welfare effects. Total surplus, defined as buyer plus seller surplus net of storage costs, is highest in the Free Market treatment. Relative to this benchmark, surplus is about 8% lower under the Anticipated Price Cap and 13% lower under Liberalization. Defensive storage can partially mitigate the inefficiency created by future rationing, whereas rent-seeking storage under liberalization compounds the misallocation associated with a currently binding price cap.

Our experiment also reveals an asymmetry between current and anticipated rationing. Participants respond strongly to rationing when it occurs in the current market, as evidenced by the growing shortage under liberalization, but respond less than equilibrium predicts when rationing is anticipated in a future market. Because this asymmetry was not anticipated when we designed the experiment, our data establish its empirical regularity but do not identify its behavioral mechanism. Two possibilities suggested by the literature reviewed below are limited strategic reasoning and preferences over the timing of uncertainty resolution.

1.1 Related literature

Our paper contributes to the literatures on price controls, hoarding and storage, and experimental markets. As classically emphasized by Hayek (1945), prices facilitate coordination

and convey information among economic actors. Price controls interfere with this mechanism and can generate resource misallocation and coordination problems (Weitzman, 1974; Yohe, 1978), as well as rent seeking and non-price competition (Wedig et al., 1989; Deacon and Sonstelie, 1989; Hatfield et al., 2016; Brekke et al., 2006; Coyne and Coyne, 2015; Guenette, 2020). Our contribution is to identify storage as an additional margin through which these distortions operate and to show that its welfare consequences depend on whether hoarding is defensive or rent seeking.

The economics of storage has a long history. Arrow et al. (1951) developed a seminal analysis of optimal inventory policies, laying the foundation for understanding the role of storage costs in market dynamics. The theoretical literature on market equilibrium under price controls is also extensive, although much of it focuses on static environments; examples include Drèze (1975), Madden (1984), and Nguyen and Whalley (1986). More recently, Hatfield et al. (2012) and Herings (2018) use matching models to study competitive equilibrium with price controls.

Our work is closest theoretically to Weitzman (1991), but the mechanism is distinct. In Weitzman’s model, price controls raise the transaction costs of acquiring goods—through queuing, searching, bribing, and related activities—and rational agents stockpile to economize on those costs. Here, hoarding arises instead from uncertainty about rationing. When a price cap creates excess demand, a buyer may be unable to purchase, and this risk alone can generate storage even when transaction costs are zero. The two theories therefore share the prediction that price controls produce hoarding but arrive at it through different channels. Our laboratory design isolates the rationing channel: participants use the same computerized trading interface across treatments and face no queues, search costs, or other transaction costs. Relatedly, Hansman et al. (2020) provide suggestive evidence that sticky prices encourage household hoarding even in the absence of the usual panic or precautionary motives.

The experimental literature on price controls has focused primarily on static markets without storage. Isaac and Plott (1981) and Smith and Williams (1981) show that price ceilings in auction markets generate shortages and reduce efficiency beyond what standard

partial-equilibrium models predict. Coursey and Smith (1983) extends this finding to posted-offer institutions, while Finley et al. (2019) documents welfare losses from price controls and rent seeking in a classroom experiment. These experiments study single-period markets in which buyers cannot transfer goods across periods. Our experiment introduces storage into an experimental market with price controls, allowing us to test whether and how buyers hoard in response to anticipated rationing. Although competitive analysis is often associated with large markets, competitive outcomes can obtain in markets with few traders and indivisible goods (Martinelli et al., 2023), making our small-market laboratory implementation suitable for testing the competitive predictions.

Our results also relate to two strands of behavioral research that may help explain the observed departures from equilibrium. First, experimental work on nonequilibrium strategic reasoning shows that participants often respond to the immediate incentives of a game without fully iterating through the responses of other participants (Crawford et al., 2013). Closely related, Dal Bó et al. (2018) find that subjects systematically underappreciate how a policy change alters other participants' behavior. In our setting, buyers may recognize the direct value of purchasing for storage while failing to internalize the fixed-point relationship between aggregate storage and the probability of future rationing. Second, models of preferences over the timing of uncertainty resolution allow choices to depend on whether uncertainty is resolved immediately or only in a later period (Kreps and Porteus, 1978). This distinction is relevant to our treatment comparison: liberalization exposes buyers to rationing in the current period, whereas APC requires them to act on an endogenous probability of rationing in the future. These theories provide possible interpretations of the observed under-response to anticipated rationing, but our experiment does not elicit the beliefs needed to distinguish limited strategic reasoning from preferences over the timing of uncertainty.

The rest of the paper is organized as follows. Section 2 presents the model and its main theoretical results. Section 3 describes the experimental design and hypotheses. Section 4 reports the experimental results. Section 5 concludes. Supplementary figures, experimental

instructions, and additional materials appear in the Appendix.

2 Conceptual framework

2.1 An economy with storage

We consider an economy with two goods, “money” and a traded good available only in indivisible units. There are $n_B \geq 1$ buyers and $n_S \geq 1$ sellers, all of whom live for two periods, $t = 1, 2$.⁵ Buyer $i \in B$ has utility

$$u_i(q_{i1}, q_{i2}, m_i) = r_i(q_{i1} + q_{i2}) + m_i,$$

where $r_i > 0$ is the buyer’s reservation value, $q_{it} \in \{0, 1\}$ is consumption in period t , and $m_i \in \mathbb{R}$ denotes final money holdings. Buyers can consume at most one unit per period but can store one unit purchased in period 1 for consumption in period 2, at a storage cost $s > 0$. We allow negative money holdings and normalize initial money endowments to zero. Buyers are ordered so that

$$r_1 \geq r_2 \geq \dots \geq r_{n_B}.$$

Seller $j \in S$ has utility

$$u_j(y_{j1}, y_{j2}, m_j) = -c_{j1}y_{j1} - c_{j2}y_{j2} + m_j,$$

where $c_{jt} > 0$ is the cost of production in period t and $y_{jt} \in \{0, 1\}$. Each seller can produce at most one unit per period. Changes in production costs across periods represent supply shocks.

⁵We present the model with a finite number of traders to match the experiment, which has eight buyers and nine sellers.

2.2 Competitive equilibrium

Suppose trade takes place freely at market prices (p_1, p_2) . Let $x_{i1} \in \{0, 1, 2\}$ and $x_{i2} \in \{0, 1\}$ denote the quantities purchased by buyer i in periods 1 and 2. A buyer who chooses $x_{i1} = 2$ consumes one unit in period 1 and stores the second for consumption in period 2, at a total monetary cost of $2p_1 + s$. Buyer i therefore solves

$$\max_{x_{i1} \in \{0, 1, 2\}, x_{i2} \in \{0, 1\}} r_i(x_{i1} + x_{i2}) - p_1 x_{i1} - p_2 x_{i2} - s \max\{x_{i1} - 1, 0\},$$

while seller j solves

$$\max_{y_{j1}, y_{j2} \in \{0, 1\}} (p_1 - c_{j1})y_{j1} + (p_2 - c_{j2})y_{j2}.$$

Market clearing requires

$$\sum_i x_{i1} = \sum_j y_{j1}, \quad \sum_i x_{i2} = \sum_j y_{j2}.$$

A competitive equilibrium consists of demand vectors (x_{i1}, x_{i2}) for each buyer, supply vectors (y_{j1}, y_{j2}) for each seller, and prices (p_1, p_2) such that buyers and sellers optimize and both markets clear.

For $t = 1, 2$, define the supply and demand correspondences

$$S_t(p) \equiv \{q \in \mathbb{Z}_+ : |\{j : c_{jt} < p\}| \leq q \leq |\{j : c_{jt} \leq p\}|\},$$

and

$$D_t(p) \equiv D(p) \equiv \{q \in \mathbb{Z}_+ : |\{i : r_i > p\}| \leq q \leq |\{i : r_i \geq p\}|\}.$$

The correspondences include all integer quantities between the strict and weak counts, reflecting arbitrary behavior by traders who are indifferent at price p . Because buyer valuations are constant across periods, demand is the same in both periods, so we write $D \equiv D_1 = D_2$.

A price p_t is myopic in period t if

$$S_t(p_t) \cap D(p_t) \neq \emptyset.$$

Thus, myopic prices are the market-clearing prices that would obtain in each period if buyers ignored the possibility of storage.

If there exists a pair of myopic prices (p_1, p_2) satisfying

$$p_1 + s \geq p_2,$$

then every competitive equilibrium price pair is myopic: storage can only weakly raise the period-1 price and lower the period-2 price relative to the myopic price sets, whereas positive storage requires $p_1 + s = p_2$. Any myopic pair satisfying this inequality supports a competitive equilibrium without storage, up to storage by indifferent buyers in the knife-edge case $p_1 + s = p_2$.

Otherwise, every pair of myopic prices satisfies $p_1 + s < p_2$. In this case every competitive equilibrium involves storage, and intertemporal arbitrage implies

$$p_2 = p_1 + s.$$

For a candidate period-1 price p_1 , define

$$\tilde{S}(p_1) = S_1(p_1) - D(p_1), \quad \tilde{D}(p_1) = D(p_1 + s) - S_2(p_1 + s), \quad (1)$$

where, for sets $A, B \subseteq \mathbb{Z}$, $A - B \equiv \{a - b : a \in A, b \in B\}$. Here, $\tilde{S}(p_1)$ is the correspondence of net period-1 supply available for storage after current consumption, while $\tilde{D}(p_1)$ is the correspondence of net period-2 demand to be met from storage, evaluated at the no-arbitrage price $p_2 = p_1 + s$. Market clearing for stored goods requires

$$\tilde{S}(p_1) \cap \tilde{D}(p_1) \neq \emptyset.$$

In this environment, the aggregate market-clearing condition is also sufficient for individual feasibility. Since $s > 0$, buyers willing to consume at $p_1 + s$ are a subset of those willing to consume at p_1 . Thus, demands can be implemented so that period-2 consumers are a subset of period-1 consumers, with some of them purchasing a second unit in period 1 for consumption in period 2. No buyer therefore consumes more than one unit per period or more than two units in total. Despite the indivisibility of the traded good, a competitive equilibrium exists in this environment by the usual supply-and-demand crossing argument.

Our experimental markets contain finitely many traders, raising the question of whether competitive equilibrium provides an appropriate benchmark for the call market participants actually play. Martinelli et al. (2023) show that, in a price-quantity strategic market game with indivisible goods and finitely many traders, every Nash equilibrium outcome with active trade is competitive if and only if at least two buyers and two sellers are willing to trade at every competitive price. Our design comfortably satisfies this condition. With storage, the competitive price band is $[70, 80]$; at every price in this interval, four buyers, with values of 110 or more, and seven sellers, with costs of 60 or less, are willing to trade. Competitive outcomes are therefore supported despite the small number of traders.

As an additional check on the price-taking benchmark, we use our experimental supply and demand schedules to calculate the largest gain available to a buyer from strategically departing from price-taking behavior while paying the resulting clearing price. In a thin single-period market, a high-value buyer could gain approximately 15 tokens—about 20% of the competitive price—by shading its bid to the relevant step of the supply schedule. In the two-period market with storage, this gain disappears: storage demand clusters at the no-arbitrage willingness to pay, thickening the market sufficiently that no individual buyer can profitably move the clearing price. Thus, the experiment tests both competitive behavior in a small market and the model’s specific predictions about storage.

2.3 Expected price controls

We now impose a price ceiling $\bar{p}_2 > 0$ in period 2. If demand exceeds supply at the ceiling, the available supply is allocated uniformly at random among buyers who demand the good. Thus, a buyer who waits until period 2 may face a continuation probability—a probability of obtaining the good—smaller than one, whereas a buyer who purchases and stores in period 1 secures consumption in advance.

We follow the usual price-taking logic of competitive equilibrium. A buyer deciding whether to store takes the continuation market as given rather than internalizing the effect of her own storage decision on that market. With finitely many buyers, this creates a discreteness at the marginal buyer: the continuation probability is smaller for those who store than for those who do not store. We accommodate this discreteness by allowing the continuation probability at the margin to lie between these two values. The resulting interval collapses under replication, as the effect of an individual buyer vanishes.

Given prices $p_1 \geq 0$ and $0 \leq p_2 \leq \bar{p}_2$ and a continuation probability ρ , buyer i solves

$$\max_{x_{i1} \in \{0,1,2\}, x_{i2} \in \{0,1\}} \{r_i(x_{i1} + \rho x_{i2}) - p_1 x_{i1} - p_2 \rho x_{i2} - s \max\{x_{i1} - 1, 0\}\},$$

while seller j chooses supply as in the competitive environment. Market clearing and rationing require

$$\begin{aligned} \sum_i x_{i1} &= \sum_j y_{j1}, & \sum_i x_{i2} &\geq \sum_j y_{j2} \text{ with equality if } p_2 < \bar{p}_2, \\ \frac{\sum_j y_{j2}}{\sum_i x_{i2} + 1} &\leq \rho \leq \frac{\sum_j y_{j2}}{\sum_i x_{i2}} \text{ if } p_2 = \bar{p}_2 \text{ and } \sum_i x_{i2} > 0, & \text{ and } \rho = 1 \text{ otherwise.} \end{aligned}$$

An equilibrium under expected price controls consists of demand vectors (x_{i1}, x_{i2}) for each buyer, supply vectors (y_{j1}, y_{j2}) for each seller, and prices (p_1, p_2) with $p_2 \leq \bar{p}_2$ and probability ρ such that buyers and sellers optimize and markets satisfy the clearing and rationing conditions.

We next characterize this equilibrium when $\rho < 1$. Given a continuation probability ρ ,

buyer i weakly prefers storing a unit in period 1 to waiting until period 2 if and only if

$$r_i - (p_1 + s) \geq \rho(r_i - \bar{p}_2).$$

For $\rho < 1$, the difference between the two sides is strictly increasing in r_i . Hence storage is positively selected on reservation values.

Suppose that the k buyers with the highest reservation values store. The remaining buyers with reservation values above the ceiling demand the good in period 2. Let

$$d_2(k) = |\{i > k : r_i \geq \bar{p}_2\}|$$

denote this residual demand and define

$$\rho_k = \min \left\{ 1, \frac{\bar{y}_2}{d_2(k)} \right\},$$

with $\rho_k = 1$ if $d_2(k) = 0$, where $\bar{y}_2 = |\{j : c_{j2} \leq \bar{p}_2\}|$ is the supply available at the ceiling. Thus, ρ_k is the continuation probability generated when the first k buyers store. Since storage removes buyers from future demand, $\rho_{k+1} \geq \rho_k$. Accordingly, if k is the marginal storer, the rationing condition in the equilibrium definition requires $\rho \in [\rho_{k-1}, \rho_k]$.

This observation allows us to construct storage demand recursively. Suppose buyers $1, \dots, k-1$ store. Buyer k takes the continuation market generated by these higher-reservation buyers as given and stores if

$$r_k - (p_1 + s) \geq \rho_{k-1}(r_k - \bar{p}_2). \tag{2}$$

As k increases, reservation values fall while the probability of obtaining the good in period 2 weakly rises. Therefore, once (2) fails, it also fails for every lower-reservation buyer. Storage demand consequently has a cutoff structure.

Equivalently, k units belong to storage demand at price p_1 whenever there exists

$$\rho \in [\rho_{k-1}, \rho_k]$$

such that

$$r_k - (p_1 + s) \geq \rho(r_k - \bar{p}_2)$$

and, if $k < n_B$,

$$r_{k+1} - (p_1 + s) \leq \rho(r_{k+1} - \bar{p}_2).$$

Let $\hat{D}(p_1)$ denote the resulting storage-demand correspondence. It is weakly decreasing in p_1 : a higher current price reduces the payoff from storage without changing the continuation payoff from waiting.

Recall from (1) that

$$\tilde{S}(p_1) = S_1(p_1) - D(p_1)$$

is the correspondence of net period-1 supply available for storage after current consumption. The correspondence $\tilde{S}$ is weakly increasing in p_1 . An equilibrium under expected price controls is therefore given by a period-1 price p_1 , storage α , and a continuation probability ρ such that

$$\alpha \in \tilde{S}(p_1) \cap \hat{D}(p_1),$$

with ρ belonging to the interval associated with the marginal storage decision, together with period-2 supply and rationing at the ceiling.

Theorem 1 *An equilibrium under expected price controls exists. Moreover, if the price ceiling binds, every equilibrium constructed above has the following properties:*

- i) Period-2 supply is weakly below the maximum competitive supply.⁶*
- ii) Rationing is effective: $\rho < 1$ whenever there is excess demand in period 2.*

⁶With indivisible units and finitely many traders, the period-2 quantity at the cap can coincide with the lowest competitive quantity, so the comparison is weak.

iii) *Storage is positively selected on reservation values: only buyers with the highest reservation values store in period 1, while lower-reservation buyers who still value consumption at the ceiling wait to purchase in the rationed period-2 market.*

iv) *The period-1 price inclusive of storage costs exceeds the period-2 price ceiling:*

$$p_1 + s > \bar{p}_2.$$

By contrast, in a competitive equilibrium with storage, $p_1 + s = p_2$.

v) *Let α and α' denote storage under rationing and in a competitive equilibrium, respectively, and let p_1 and p'_1 denote the corresponding period-1 prices. Then*

$$(p_1 - p'_1)(\alpha - \alpha') \geq 0.$$

Thus, the period-1 price and the quantity stored cannot move in opposite directions.

Proof. Suppose first that some competitive equilibrium has a period-2 price $p_2 \leq \bar{p}_2$. The price ceiling does not constrain that equilibrium, which is therefore also an equilibrium under expected price controls.

Now suppose that $\bar{p}_2$ lies below every competitive period-2 price. The ceiling must then bind. Otherwise period 2 would clear at some price $p_2 \leq \bar{p}_2$, and together with optimization and period-1 market clearing this would constitute a competitive equilibrium with a period-2 price no greater than the ceiling, a contradiction.

Period-2 supply at the ceiling is weakly below the maximum competitive period-2 supply, proving (i). Whenever period-2 demand at the ceiling exceeds available supply, the rationing condition implies $\rho < 1$, proving (ii).

Part (iii) was established in the construction of storage demand above: for $\rho < 1$ the gain from storing rather than waiting is strictly increasing in the reservation value, and the recursive construction generates an upper interval of storers.

For (iv), suppose, to the contrary, that

$$p_1 + s \leq \bar{p}_2.$$

Then any buyer who obtains positive surplus from consuming in period 2 weakly prefers securing the unit through storage at the weakly lower acquisition cost, and strictly prefers doing so whenever rationing is effective. There can therefore be no excess demand in period 2 consistent with a binding ceiling. Hence

$$p_1 + s > \bar{p}_2.$$

For (v), let α denote storage in an equilibrium under expected price controls and α' storage in a competitive equilibrium, with corresponding period-1 prices p_1 and p'_1 . By period-1 market clearing,

$$\alpha \in \tilde{S}(p_1), \quad \alpha' \in \tilde{S}(p'_1).$$

The correspondence $\tilde{S}$ is weakly increasing: a higher period-1 price weakly increases supply and weakly reduces current consumption. Its values are ordered, so

$$p_1 > p'_1 \implies \alpha \geq \alpha'.$$

Conversely, $\alpha > \alpha'$ and $p_1 < p'_1$ would contradict this monotonicity. The corresponding argument in the opposite direction gives

$$(p_1 - p'_1)(\alpha - \alpha') \geq 0.$$

It remains to establish existence when the ceiling binds. The storage-demand correspondence $\hat{D}(p_1)$ is weakly decreasing in p_1 , while the net supply correspondence $\tilde{S}(p_1)$ is weakly increasing. At sufficiently low period-1 prices, net supply available for storage is no greater than storage demand, whereas at sufficiently high prices storage demand vanishes while net

supply is nonnegative. Because both correspondences include the intervals generated by indifference and by the finite-market change in the continuation probability at the marginal buyer, they intersect when their ordering reverses. Hence there exists a period-1 price p_1 and a quantity α such that

$$\alpha \in \tilde{S}(p_1) \cap \hat{D}(p_1).$$

The associated continuation probability ρ and period-2 rationing complete the equilibrium. This establishes existence. ■




Figure 1: Discrete market for storage under an anticipated price cap. Rationing raises storage demand relative to the competitive correspondence $\tilde{D}$, increasing equilibrium storage from $\alpha_c = 2$ to $\alpha_r = 3$.

Figure 1 illustrates the discrete storage market for the schedules used in our experiment. Under the anticipated price cap, rationing raises storage demand relative to competitive storage demand $\tilde{D}$. Because buyers anticipate a positive probability of being rationed in period 2, they are willing to pay more for an additional unit in period 1 and store it for future consumption. For the experimental parameters, equilibrium storage rises from $\alpha_c = 2$ to $\alpha_r = 3$. Storage mitigates the future shortage but does not eliminate it.

2.4 Expected liberalization

We now reverse the timing of the price distortion. A price ceiling $\bar{p}_1 > 0$ applies in period 1 and is removed in period 2, when the price is freely determined. Buyers may purchase an additional unit in period 1 and store it for consumption in period 2.

Rationing in period 1 is symmetric across requested units. Each buyer requests one unit for current consumption, or two units—one for current consumption and one for storage. At any price below the ceiling the market clears and every request is served. If the number of requested units exceeds the supply available at the ceiling, the available units are allocated uniformly at random among all requested units, with no distinction between first and second requests. A buyer who requests two units stores only if she obtains both; if she obtains a single unit, she consumes it in period 1 and may still purchase in period 2. Thus, both the number of units stored and the identities of the storers are determined by the allocation lottery.

The identities of successful storers matter in a finite market. A buyer who successfully stores does not return to purchase in period 2. Since storage units are allocated randomly rather than according to reservation values, different realizations of the storage lottery generally leave different compositions of residual period-2 demand and may therefore generate different competitive prices after liberalization. A buyer contemplating storage must consequently evaluate a distribution of possible period-2 prices rather than a single future price.

As under expected price controls, we follow the price-taking logic of competitive equilibrium. The continuation market is now summarized by three objects, which we call a continuation assessment: the probability ρ that a requested unit is obtained, the probability $\pi \leq \rho$ that a buyer requesting two units obtains both, and the distribution F of the period-2 market-clearing price. A buyer deciding whether to request a second unit takes the assessment as given and does not internalize the effect of her own request on it. With finitely many buyers, adding the marginal applicant changes the assessment. We accommodate this discreteness by allowing the common assessment at the margin to be a mixture of the assessments generated immediately before and after the marginal buyer enters the request pool;

expected payoffs under a mixture of assessments are the corresponding mixtures of expected payoffs.

An equilibrium under expected liberalization consists of period-1 requests and supplies, a period-1 price $p_1 \leq \bar{p}_1$, and a common continuation assessment such that sellers optimize, each buyer's requests are optimal taking the price and the assessment as given, the period-1 market satisfies the clearing and rationing conditions described above, and the period-2 market clears competitively for every realization of the allocation lottery. If some buyer requests a second unit, let the marginal storer be the requester with the lowest reservation value. The common assessment must be a mixture of the assessment induced by the equilibrium requests and the assessment induced when the marginal storer requests a single unit, holding all other requests fixed. If no buyer requests a second unit, the assessment is the one induced by single-unit requests.

We next characterize the equilibrium when the period-1 ceiling binds, so that $p_1 = \bar{p}_1$ and requested units exceed the available supply. Given an assessment (ρ, π, F) , a buyer who requests a single unit obtains it with probability ρ and retains the option to purchase in period 2, for an expected payoff of

$$\rho(r_i - \bar{p}_1) + \mathbb{E}_F [(r_i - p_2)^+], \quad (z)^+ \equiv \max\{z, 0\}.$$

If she requests two units, she obtains both with probability π , in which case she consumes one unit, stores the other, and does not purchase in period 2; she obtains exactly one with probability $2(\rho - \pi)$, in which case she consumes it; and she retains the period-2 purchase option whenever she does not store. Her expected payoff is therefore

$$\begin{aligned} \pi[2(r_i - \bar{p}_1) - s] + 2(\rho - \pi)(r_i - \bar{p}_1) + (1 - \pi) \mathbb{E}_F [(r_i - p_2)^+] \\ = 2\rho(r_i - \bar{p}_1) - \pi s + (1 - \pi) \mathbb{E}_F [(r_i - p_2)^+]. \end{aligned}$$

The gain from the second request is the difference between the two expected payoffs,

$$g(r_i) = \rho(r_i - \bar{p}_1) - \pi \left(s + \mathbb{E}_F \left[(r_i - p_2)^+ \right] \right). \quad (3)$$

A second request thus has two effects: it raises the probability of obtaining a unit for current consumption, and it creates the possibility of storage, which costs s and replaces the period-2 purchase option. Buyer i weakly prefers to request a second unit if and only if $g(r_i) \geq 0$.

For a given assessment, the gain is increasing in the reservation value. For $r' > r$,

$$\mathbb{E}_F \left[(r' - p_2)^+ - (r - p_2)^+ \right] \leq r' - r,$$

and hence

$$g(r') - g(r) \geq (\rho - \pi)(r' - r) \geq 0,$$

since $\pi \leq \rho$. Thus, second-unit requests have a cutoff structure for any given assessment.

Suppose buyers $1, \dots, k$ request a second unit, and let d_1 denote the number of buyers requesting a unit for current consumption at the ceiling, so that $d_1 + k$ units are requested in total. With

$$\bar{y}_1 = |\{j : c_{j1} \leq \bar{p}_1\}|$$

units supplied at the ceiling, the uniform lottery yields

$$\rho_k = \frac{\bar{y}_1}{d_1 + k}, \quad \pi_k = \rho_k \frac{\bar{y}_1 - 1}{d_1 + k - 1}$$

when $d_1 + k > \bar{y}_1$, and $\rho_k = \pi_k = 1$ when all requests can be served. The lottery over successful storers induces a distribution F_k of period-2 competitive market-clearing prices. Let g_k denote the gain function (3) evaluated at the assessment (ρ_k, π_k, F_k) .

In contrast with expected price controls, the induced assessments need not be monotone in k : an additional request removes a high-reservation buyer from residual period-2 demand in the realizations in which she stores, but it also dilutes the probability that every other applicant completes a pair, and the two effects move the period-2 price in opposite directions.

If buyer k is the marginal storer, the admissible assessments are the mixtures of the assessments generated immediately before and after her request, with associated gain functions

$$g^\lambda = \lambda g_{k-1} + (1 - \lambda)g_k, \quad \lambda \in [0, 1].$$

Each g^λ is weakly increasing in r , being a convex combination of weakly increasing functions. The first k buyers constitute equilibrium storage demand whenever there exists $\lambda \in [0, 1]$ such that

$$g^\lambda(r_k) \geq 0$$

and, if $k < n_B$,

$$g^\lambda(r_{k+1}) \leq 0.$$

These two inequalities imply that buyers $1, \dots, k$ weakly prefer to request a second unit while buyers $k + 1, \dots, n_B$ weakly prefer not to. The mixture interval is the counterpart of the interval of rationing probabilities associated with the marginal buyer under expected price controls.

The finite-market interval disappears under replication. As the economy is replicated, the composition of successful storers converges to its aggregate composition, so the distribution of the future market-clearing price collapses to a single price p_2^* , and the effect of the marginal applicant on the assessment vanishes. The mixture interval then collapses to a singleton, and the construction reduces to an exact cutoff rule.

Theorem 2 *An equilibrium under expected liberalization exists. Moreover, if the period-1 price ceiling binds:*

- i) Period-1 supply is weakly below competitive supply, and rationing occurs whenever demand at the ceiling exceeds available supply.*
- ii) Storage requests are positively selected on reservation values: only buyers with the highest reservation values request storage, up to possible indifference at the cutoff.*

iii) *If positive storage occurs, residual period-2 demand is weakly below myopic period-2 demand in every realization. Consequently, every realized period-2 market-clearing price belongs to the market-clearing price correspondence associated with demand weakly below myopic demand.*

Proof. If the period-1 ceiling does not bind, the corresponding competitive equilibrium is also an equilibrium under expected liberalization. Suppose henceforth that the ceiling binds.

Part (i) follows directly from the price ceiling and the rationing rule. At $p_1 = \bar{p}_1$, sellers supply only units whose production cost does not exceed the ceiling, so period-1 supply is weakly below competitive supply. If the number of requested units exceeds this supply, the uniform lottery rations every request with the same probability $\rho < 1$.

For part (ii), it was shown above that the gain from a second request is weakly increasing in r for any admissible assessment. Hence, if a buyer with reservation value r_i requests a second unit, every buyer with a higher reservation value weakly prefers to do so, and second-unit requests form an upper interval of the reservation-value ranking, up to indifference at the cutoff.

For part (iii), fix any realization of period-1 rationing with positive storage. Every stored unit weakly reduces period-2 demand: if the successful storer would otherwise purchase at the realized period-2 price, her demand is removed, while otherwise demand is unchanged. Relative to the myopic period-2 market, supply is unchanged and demand is therefore weakly lower at every price. The conclusion holds realization by realization and hence for the entire distribution of period-2 prices.

It remains to establish existence when the ceiling binds. If

$$g_0(r_1) \leq 0,$$

zero second-unit requests constitute an equilibrium: the assessment is the one induced by single-unit requests, and since g_0 is weakly increasing in r , no buyer gains from requesting a second unit. In this case period-1 rationing does not affect period-2 demand, because buyers who are rationed in period 1 may still purchase for period-2 consumption, so F_0 is

the distribution of competitive prices generated by myopic period-2 demand.

Suppose instead that $g_0(r_1) > 0$. Walk down the reservation-value ranking, adding one buyer at a time to the request pool, and stop the first time the next buyer declines to join the pool generated by those above her. Formally, let k^* be the smallest index $k \in \{1, \dots, n_B\}$ with

$$g_k(r_{k+1}) \leq 0,$$

and set $k^* = n_B$ if no such index exists. Because the walk did not stop earlier, every $k < k^*$ satisfies $g_k(r_{k+1}) > 0$; restated with the index shifted by one, and using the hypothesis $g_0(r_1) > 0$ for $k = 1$,

$$g_{k-1}(r_k) > 0 \quad \text{for every } k \leq k^* :$$

each buyer along the walk strictly preferred to join the pool she faced. In particular, $g_{k^*-1}(r_{k^*}) > 0$.

If $k^* = n_B$, all buyers request two units. Taking $\lambda = 1$, so that the common assessment is the one generated without the marginal request, monotonicity in r yields

$$g^\lambda(r_i) \geq g^\lambda(r_{n_B}) = g_{n_B-1}(r_{n_B}) > 0$$

for every buyer i , so all requests are optimal.

Suppose then that $k^* < n_B$, so that

$$g_{k^*}(r_{k^*+1}) \leq 0.$$

The walk has found a seam: buyer k^* strictly preferred to join under the assessment generated without her request, while buyer $k^* + 1$ declines under the assessment generated with it. The two assessments disagree about the margin, and mixing them produces the required indifference. For $\lambda \in [0, 1]$, let $g^\lambda = \lambda g_{k^*-1} + (1 - \lambda) g_{k^*}$, so that $\lambda = 1$ is the assessment without the marginal request and $\lambda = 0$ the one with it, and track the gains of the two

buyers at the seam:

$$A(\lambda) = g^\lambda(r_{k^*}), \quad B(\lambda) = g^\lambda(r_{k^*+1}),$$

the gain of the marginal buyer and of the first excluded buyer. Both functions are affine in λ , with

$$A(1) = g_{k^*-1}(r_{k^*}) > 0, \quad B(0) = g_{k^*}(r_{k^*+1}) \leq 0,$$

and $A(\lambda) \geq B(\lambda)$ for every λ , because each g_k is weakly increasing in r and $r_{k^*} \geq r_{k^*+1}$. If $B(1) \leq 0$, set $\lambda^* = 1$; then $A(\lambda^*) = A(1) > 0$. Otherwise $B(0) \leq 0 < B(1)$, and by affinity there exists λ^* with $B(\lambda^*) = 0$, whence $A(\lambda^*) \geq B(\lambda^*) = 0$. In either case,

$$A(\lambda^*) \geq 0 \quad \text{and} \quad B(\lambda^*) \leq 0 :$$

under the assessment g^{λ^*} the marginal buyer weakly prefers to request two units and the first excluded buyer weakly prefers not to, and by monotonicity in r the same holds for every buyer above and below them, respectively. Together with the uniform allocation lottery and competitive period-2 clearing for every realization, this constitutes an equilibrium under expected liberalization. Hence an equilibrium exists.⁷ ■

Figure 2 illustrates the discrete period-1 market in the Liberalization treatment. The price ceiling $\bar{p}_1$ lies below the free market-clearing price, so requested units exceed supply and buyers are rationed. Buyers with sufficiently high reservation values have an incentive to request a second unit and store it in anticipation of the liberalized period-2 market. Because storage requires obtaining both requested units, and $\pi \leq \rho$, the effective storage cost at the margin is $(\pi/\rho)s$, and the request threshold $\bar{p}_1 + (\pi/\rho)s$ lies below the storage threshold $\bar{p}_1 + s$ that would apply if storage were certain.

⁷We do not claim uniqueness in general: a lower cutoff dilutes the probability that each applicant completes a pair, which both cheapens a second request and returns high-reservation buyers to the period-2 market, so requests are strategic complements at the margin. For the parameters of our experiment, enumerating all candidate cutoffs shows that the equilibrium is unique: exactly the six buyers with valuations of at least 70 request a second unit.




Figure 2: Discrete period-1 market under expected liberalization. The price ceiling $\bar{p}_1 = 50$ binds below the free market-clearing price, so requested units exceed supply and every request faces the same allocation lottery. Buyers with reservation values above the request threshold $\bar{p}_1 + (\pi/\rho)s$ request a second unit in anticipation of the liberalized period-2 market.

3 Experimental design

3.1 Hypotheses

The conceptual framework yields a set of predictions for prices, storage, and shortages across our three experimental environments. The parameter values used in the experiment satisfy the conditions of the theoretical results above, so the corresponding comparative predictions apply directly.

Hypothesis 1 (Free market) *In a competitive market with storage, intertemporal arbitrage implies*

$$p_1 + s = p_2.$$

Buyers are therefore indifferent between purchasing an additional unit in period 1 and storing it for period 2 and purchasing the unit directly in period 2. Markets clear without rationing.

Hypothesis 2 (Expected price controls) *Suppose that a binding price ceiling $\bar{p}_2$ is anticipated in period 2. Then:*

- a) *Buyers with the highest reservation values store in period 1, while buyers with relatively lower reservation values who still demand the good wait to purchase in the rationed period-2 market.*
- b) *The period-1 price inclusive of storage costs strictly exceeds the period-2 price ceiling:*

$$p_1 + s > \bar{p}_2.$$

- c) *The period-1 price is weakly higher than the competitive period-1 price.*
- d) *The quantity stored is weakly greater than in the competitive equilibrium.*

Hypothesis 3 (Expected liberalization) *Suppose instead a binding price ceiling $\bar{p}_1$ is imposed in period 1 and the market is liberalized in period 2. Then:*

- a) *Period-1 demand exceeds supply at the price ceiling, generating a shortage and effective rationing.*
- b) *Buyers with sufficiently high reservation values request a second unit for storage. Because all requested units enter the same lottery, a second request both raises the chance of obtaining a unit for current consumption and creates the possibility of storage. A buyer weakly prefers to request a second unit if and only if*

$$\rho(r_i - \bar{p}_1) \geq \pi(s + \mathbb{E}[(r_i - p_2)^+]),$$

where ρ is the probability that a requested unit is obtained and $\pi \leq \rho$ is the probability that both requested units are obtained, as in equation (3). The gain is increasing in r_i , so storage requests have a cutoff structure. For the experimental parameters the cutoff is $\bar{p}_1 + (\pi/\rho)s \approx 53$: the six buyers with valuations of at least 70 request two units, while buyers with valuations of 50 do not.

- c) *If there is storage, residual period-2 demand is weakly below myopic period-2 demand in every realization, so every realized period-2 price is weakly below the myopic price range.*

For the experimental parameters, the predicted price distribution concentrates on $[90, 120]$, weakly below the myopic range $[110, 120]$.

3.2 Experiment

The experiment consists of a sequence of two-period markets for a durable good. In period 1, buyers can purchase up to two units, consuming one immediately and storing the second at a cost for consumption in period 2. A buyer who successfully acquires a unit for storage in period 1 does not participate in the period-2 market.

We use a between-subjects design, assigning each session to one treatment. Each participant begins a round with a budget of 200 laboratory currency units (LCU). Buyer valuations are randomly assigned across rounds so that participants have the same expected earnings over the course of the experiment.

Sessions consist of either 8 or 16 participants, who complete 10 rounds, each containing two trading periods. Each market has eight human buyers and nine automated (“robot”) sellers programmed to follow equilibrium strategies. In eight-participant sessions, the same eight buyers form the market throughout the session. In sixteen-participant sessions, participants are randomly rematched at the beginning of each round into two groups of eight and interact only with participants from the same session. We refer to each group-round pair as a market and, for statistical inference, treat the session as the independent matching group (Section 4).

The use of automated sellers allows us to isolate demand-side hoarding by holding seller behavior fixed and reducing noise from strategic interaction on both sides of the market. Automated sellers have been used previously in experimental markets (Cason and Friedman, 1997). In our setting, they also rule out seller collusion, strategic supply withdrawal, and seller-side learning, so departures from the theoretical predictions can be attributed more directly to buyer behavior.

As discussed in Section 2, our design satisfies the condition in Martinelli et al. (2023) under which all Nash equilibrium outcomes with active trade are competitive: at least two

buyers and two sellers are willing to trade at every competitive price. Competitive outcomes therefore provide an appropriate benchmark for the free-market periods, while periods with a binding price ceiling are evaluated against the corresponding rationing equilibrium derived above. The experiment thus jointly tests competitive behavior in a small market and the model’s specific predictions concerning storage, prices, and rationing.

The use of automated sellers comes at the cost of abstracting from potentially important supply-side responses to price controls. Our design cannot address quality degradation, strategic supply withdrawal, pass-through, or joint learning by buyers and sellers, nor how these responses might interact with hoarding. These questions are left for future experiments. Detailed instructions and experimental materials are provided in the Appendix.

3.3 Market institution

We use a single call market (SCM) institution. SCM asks each buyer i to submit their maximum acceptable price, bid b_i , and to each seller to submit their minimum acceptable sale price, ask a_j . The market is cleared at a single equilibrium price p^* . This particular market institution reduces behavioral noise, guarantees a unique equilibrium price, and human subjects seem to behave more efficiently when interacting with automated agents (Cason and Friedman, 1997).

3.4 Implementation of the model

We utilize three different treatments for this experiment. In all treatments, buyers’ valuations remain constant across both periods, while sellers’ production costs increase fourfold in the second period ($c_2(i) = 4 \times c_1(i)$). The treatments differ in whether and when a price ceiling is imposed.

1. **Free Market:** This treatment serves as the baseline. There is no price cap in either period. Sellers’ costs are four times greater in the second period.

Demand: $r(i) = [50, 50, 70, 90, 110, 140, 140, 200]$

Supply: $c_1(i) = [5, 20, 20, 30, 50, 55, 60, 80, 100]$ and
 $c_2(i) = [20, 80, 80, 120, 200, 220, 240, 320, 400]$.

In the myopic equilibrium, the quantity traded in the first period is $q_1 = 6$, with the corresponding price falling within the range of $[55, 60]$. Similarly, the equilibrium quantity traded in the second period is $q_2 = 3$, and the corresponding price is in the range of $p_2 \in [110, 120]$.

Notice that $p_2 > p_1 + s$ indicates an incentive to store goods in period 1. As a result, the predicted competitive equilibrium leads to the following outcomes:

- Equilibrium probability: $\rho = 1$ (no rationing in equilibrium).
- Demand for storage: $\tilde{D} = [50, 70, 70, 70, 80, 110]$
- Supply for storage: $\tilde{S} = [60, 70, 80, 90, 100, 110, 140, 140, 200]$
- In the competitive equilibrium, the quantity traded in the first period can be either $q_1^* = 7$ or $q_1^* = 8$, with the corresponding price falling within the $[70, 80]$ range. $\alpha^* = 2$ units are put in storage. For the second period, the equilibrium quantity traded is $q_2^* = 3$, and the corresponding price is in the range of $[80, 90]$.

[Figure 3 about here]

2. **Anticipated Price Cap:** In this treatment, we introduce a price cap regulation in the second period. Buyers' valuations remain constant across both periods, and sellers' costs are four times greater in the second period, as in the Free Market treatment. In addition, a price ceiling is imposed in the second period.

Demand: $r(i) = [50, 50, 70, 90, 110, 140, 140, 200]$

Supply: $c_1(i) = [5, 20, 20, 30, 50, 55, 60, 80, 100]$ and
 $c_2(i) = [20, 80, 80, 120, 200, 220, 240, 320, 400]$.

The price ceiling, $\bar{p}_2$, is set at 50 in period 2. After the supply shock, only one seller has a cost below the ceiling ($c_{j2} = 20 \leq 50$), so $S_2(\bar{p}_2) = 1$ while $D(\bar{p}_2) = 8$. This

severe shortage triggers a strong demand for storage with a cost of $s = 10$. As a result, the storage demand under rationing, listing each buyer's willingness to pay for a stored unit at the continuation probabilities of the recursive construction of Section 2, is $\hat{D} = [171.25, 117.14, 115, 88, 70, 53.33, 40]$, while the supply for storage is $\tilde{S} = [60, 70, 80, 90, 100, 110, 140, 140, 200]$. Therefore, the predicted market equilibrium is as follows:

- Equilibrium probability of acquiring a good in period 2 is $\rho \in [0.17, 0.20)$.
- In the rationing equilibrium, the quantity traded in the first period is $q_1^* = 8$, among them $\alpha^* = 3$ are put in storage, with the corresponding price p_1^* falling within the range $[88, 90]$. In the second period, the traded quantity is $q_2^* = 1$ (only the lowest-cost seller can sell below the ceiling), with the corresponding price $p_2^* = \bar{p}_2 = 50$.

[Figure 4 about here]

3. **Liberalization:** In this treatment, we introduce a price cap regulation in the first period that is removed in the second period. Buyers' valuations remain constant across both periods, and sellers' costs are four times greater in the second period, as in the other treatments.

Demand: $r(i) = [50, 50, 70, 90, 110, 140, 140, 200]$

Supply: $c_1(i) = [5, 20, 20, 30, 50, 55, 60, 80, 100]$ and

$c_2(i) = [20, 80, 80, 120, 200, 220, 240, 320, 400]$.

The price ceiling, $\bar{p}_1$, is set at 50 in period 1, and is binding, resulting in a shortage during that period. The cost of storage is set at $s = 10$. The predicted market equilibrium is as follows.

- The supply at the ceiling in period 1 is $S_1(\bar{p}_1) = 5$ (five sellers have costs ≤ 50), while the demand is $D(\bar{p}_1) = 8$. The equilibrium probability of acquiring a good

in period 1 is $\rho = 5/14$, as 6 buyers with valuations ≥ 70 demand 2 widgets and 2 buyers with valuations of 50 demand 1 widget each.

- In the rationing equilibrium, the quantity traded in the first period is $q_1^* = 5$, with the corresponding price $p_1 = 50$, the price ceiling.
- Because the storage decision is made within the rationed first period, the number of units placed in storage is itself a lottery outcome. All requested units enter a single uniform lottery, so a buyer stores only if both of her requested units are successful; a buyer who obtains a single unit consumes it in period 1. For each of the six two-unit bidders, the probability of obtaining both units is $(5/14)(4/13) = 10/91 \approx 0.11$, so expected storage is $\mathbb{E}[\alpha^*] = 6 \times 10/91 \approx 0.66$ units. The realized count is $\alpha^* = 0$ with probability 0.42, $\alpha^* = 1$ with probability 0.51, and $\alpha^* = 2$ with probability 0.07; storage of three or more units is impossible, since the five available units contain at most two complete pairs.⁸
- In the second period, the traded quantity is $q_2^* = 3$ for every realization of the storage lottery. The market-clearing price is $p_2^* \in [110, 120]$ (the myopic range) when no units are stored, and lies in $[80, 120]$ otherwise, depending on the identities of the successful storers. Averaging over the storage lottery, the predicted price distribution concentrates on $[90, 120]$, which carries 95% of the probability mass, with a mean of approximately 109 (taking the midpoint of each realized price range).

[Figure 5 about here]

We administer this experiment using an online platform developed in oTree (Chen et al., 2016). At George Mason University, we organized 23 sessions, each involving 8 or 16 students, all acting as buyers. In one Liberalization session the application failed to record market

⁸This contrasts with the Free Market ($\alpha^* = 2$) and Anticipated Price Cap ($\alpha^* = 3$) treatments, where the storage decision is made in a free period and storage is deterministic. Here, because the ceiling binds in the first period, realized storage depends on the period-1 rationing draw.

outcomes (prices, quantities, and bids are missing in nine of ten rounds), and that session is excluded; no other session was dropped.⁹ Our final sample consists of 288 participants in 22 sessions: 88 participants in 6 Free Market sessions, 112 in 9 Anticipated Price Cap sessions, and 88 in 7 Liberalization sessions. Of these, 164 are male, 117 are female, 4 identify as non-binary, and 3 chose not to reveal their gender. About 60% of participants are enrolled in STEM programs. Table 1 details the results of our post-experiment demographic survey. Participant composition is balanced across the three treatments: the gender distribution does not differ significantly between treatments (χ^2 test, $p = 0.94$), and the share of STEM students is comparable across treatments.

[Table 1 about here]

4 Experimental results

4.1 Overview

We organize the results around three questions. Does the free market satisfy the intertemporal arbitrage benchmark? Do buyers respond to an anticipated price cap by storing defensively? And how does behavior change when a binding cap is later removed? The evidence supports the benchmark price and shortage predictions, but it also reveals a sharp asymmetry: buyers respond weakly to anticipated rationing and strongly to rationing they experience in the current market.

Unless otherwise noted, formal inference treats the session as the independent unit of observation. Participants interact repeatedly with others in the same session, so neither buyer-rounds nor market-rounds are independent observations. There are six sessions in the Free Market treatment, nine in the Anticipated Price Cap treatment, and seven in Liberalization. We report market-level means and standard errors to describe the data, and

⁹The replication package documents this rule and its history. Fourteen sessions had 16 participants and eight had 8.

explicitly label the occasional market-level test. With so few independent units, the session-level nonparametric tests are necessarily conservative: with six sessions, for example, the smallest attainable two-sided Wilcoxon p -value is 0.031. Because treatments were assigned at the session level, we also report exact permutation tests on the session means: randomization inference over all reassignments of sessions to treatments for between-treatment comparisons, and sign-flip tests for comparisons with a theoretical benchmark. These use the magnitudes rather than the ranks of the session means and lead to the same conclusions as the rank tests. Appendix A.3 reports both sets of tests and shows that using the market groups (eleven, fourteen, and eleven per treatment) as the unit, as in an earlier version of the paper, leads to the same qualitative conclusions. Unless otherwise noted, results use the last five rounds (rounds 6–10), after participants have gained experience with the institution and incentives. The experiment was not preregistered; Appendix A.2 replicates the main results using all ten rounds.

4.1.1 Free Market

H1: *In a competitive market with storage, $p_1 + s = p_2$: buyers are indifferent between storing and purchasing in period 2, and markets clear without rationing.*

Results: The Free Market treatment closely matches the competitive price and shortage predictions. The first-period price averages 73.67, within the predicted range of [70, 80], and the second-period price averages 88.15, within the predicted range of [80, 90]. We observe no rationing. The no-arbitrage condition $p_1 + s = p_2$, however, holds only approximately. Table 4 reports an average excess gap of $p_2 - p_1 - s = 4.47$. This gap is different from zero in a market-level t -test ($t = 2.32$, $p = 0.024$, $N = 55$), but it is positive in only five of the six independent sessions and the corresponding session-level Wilcoxon signed-rank test gives $p = 0.219$. Thus, the session-level evidence does not reject the arbitrage benchmark, although the positive point estimate suggests that first-period bids did not fully eliminate the observed price spread.

[Table 4 about here]

[Figure 10 about here]

4.1.2 Anticipated Price Cap

H2a: *Buyers with the highest reservation values store in period 1, while buyers with relatively lower reservation values who still demand the good wait to purchase in the rationed period-2 market at the price ceiling $\bar{p}_2$.*

Results: We have sufficient evidence to confirm this hypothesis. There is a statistical difference in valuations between participants who chose to buy two goods in the first period and those who did not: storers have a median valuation of 110 versus 90 for non-storers (Mann–Whitney test on buyer-rounds, $p < 0.001$), and the mean valuation of storers exceeds that of non-storers in every session (by 10 to 28 tokens; Wilcoxon signed-rank test on the nine within-session differences, $p = 0.004$). Buyers with higher values were thus more inclined to buy two goods. Note, however, that this concerns who obtains two units; the relationship between valuation and the quantity buyers bid for is much weaker in every treatment (see Table 2 and the bidding-strategy discussion below).

H2b: *The period-1 price inclusive of storage costs strictly exceeds the period-2 price ceiling: $p_1 + s > \bar{p}_2$.*

Results: Our data support this hypothesis. With the introduction of price cap regulation at $t = 2$, the average first-period price (73.0) was significantly higher than the price ceiling of 50 set for the subsequent period (one-sample Wilcoxon signed-rank test on session averages, $p = 0.004$, $N = 9$ sessions; the exact sign-flip permutation test gives the same value; every session average exceeds 58). The ordering the theory predicts therefore holds: $p_1 + s = 83.0 > \bar{p}_2 = 50$. The level of the first-period price, however, falls short of the rationing-equilibrium range; we return to this under H2c.

H2c: *The period-1 price is weakly higher than the competitive period-1 price.*

Results: In the Anticipated Price Cap treatment, the first-period price averages 73.00, which is above the myopic equilibrium of [55, 60] but below the rationing equilibrium prediction of [88, 90]. The price in the first period exceeds what would prevail without storage

incentives. However, Table 5 shows that the difference in first-period prices between the Free Market and the Anticipated Price Cap treatments is not statistically significant.

H2d: *The quantity stored is weakly greater than in the competitive equilibrium.*

Results: In the Anticipated Price Cap treatment, participants stored 2.31 units per market, which exceeds the competitive equilibrium of 2 units but falls short of the rationing equilibrium prediction of 3 units. The shortfall relative to the rationing prediction is statistically significant (one-sample Wilcoxon signed-rank test on session averages against 3: $p = 0.004$, $N = 9$ sessions, the same value under the exact permutation test; every session stores fewer than 3 units on average), while storage is not significantly above the competitive level of 2 (Wilcoxon $p = 0.109$; permutation $p = 0.094$). This is consistent with the broader pattern observed throughout our experiment: participants store in the right direction but less aggressively than predicted, as if they underestimate the severity of future rationing.

4.1.3 Liberalization

H3a: *Period-1 demand exceeds supply at the price ceiling, generating a shortage and effective rationing.*

Results: We confirm this hypothesis. Shortages are near zero in all periods without a price cap. In the Liberalization treatment, where a price cap binds in the first period, the average shortage rate is 25%. As Figure 14 shows, rates between 50% and 60% are the most frequent outcome, in line with the theory. The shortage rate rises from 12% in round 1 to 34% in round 10. This increase appears across independent sessions: the shortage rate rises over rounds in every one of the seven liberalization sessions, the median session-specific trend is 1.8 percentage points per round, and a Wilcoxon signed-rank test rejects a zero median trend ($p = 0.016$, $N = 7$ sessions). Because the measure counts unsatisfied demand among bids at or above the transaction price, the result means that more demand becomes effective at the controlled price over time. Participants appear to learn how to position their bids for regulated units before liberalization.

H3b: *Storage requests have a cutoff structure: only buyers with reservation values above*

the request threshold $\bar{p}_1 + (\pi/\rho)s \approx 53$ request a second unit. For the experimental parameters, buyers with valuations of at least 70 request two units, while buyers with valuations of 50 do not.

Results: The evidence is directionally consistent with this hypothesis but weak. In the last 5 rounds, 59% of the bids submitted by buyers with valuations above the request threshold (valuations of 70 or higher) request two units, against 52% for buyers with the lowest valuation (50). The difference is not statistically significant (Pearson’s chi-squared $p = 0.22$; a Wilcoxon signed-rank test on the within-session difference in attempt rates gives $p = 0.20$, $N = 7$ sessions), and Table 2 shows the same picture for the full valuation distribution. Buyers above the threshold do request a second unit somewhat more often, as the theory predicts, but far from all of them do, and low-valuation buyers request almost as often. The theory itself, however, predicts only weak selection on this margin: because all requested units enter the same lottery, a second request is costly only if it succeeds, and for a valuation-50 buyer the expected loss from an out-of-equilibrium request is only about one token. The near-uniform two-unit bidding observed across valuations is therefore less anomalous than a pure storage motive would suggest.

H3c: *If there is storage, every realized period-2 price is weakly below the myopic price range: the predicted price distribution concentrates on $[90, 120]$, with a mean of approximately 109, weakly below the myopic range $[110, 120]$.*

Results: The data support this hypothesis. The second-period price averages 109.04, essentially coinciding with the mean of the predicted price distribution (approximately 109) and slightly below the lower bound of the myopic price range $[110, 120]$. At the level of the seven independent sessions, four have an average price below 110 and three have an average above 110; a one-sided Wilcoxon signed-rank test does not reject equality with the myopic benchmark ($p = 0.766$). This pattern is consistent with the prediction rather than evidence against it: because equilibrium storage is rare—the storage lottery yields no stored units in about 42% of markets—the predicted price distribution places substantial mass on the myopic range itself, and the predicted departure from the myopic benchmark is small. The

clearer result is the change across periods: the average price rises from the first-period cap of 50 to 109.04 following liberalization. Buyers who stored at the regulated price paid a total of 60, including the storage cost, for goods that would cost approximately 109 in the free second-period market.

4.2 Bidding strategy

The remaining subsections examine the behavior behind these aggregate results: bidding, market prices, consumption and hoarding, and shortage and efficiency.

We start by examining the bidding strategies adopted by the participants in each treatment.¹⁰ Figures 6 and 7 illustrate the average price offered in each period based on the valuation assigned during the experiment. From these graphs, it is evident that participants generally offer prices below their valuation, as expected. The prices offered in the initial five rounds are slightly higher than those in the final five rounds, indicating adjustments. However, there are no significant differences between the outcomes of the first and last 5 rounds.

Furthermore, although bid prices tend to gravitate near the equilibrium, there is a tendency to overbid in the first period and underbid in the second, especially in the Free Market and Anticipated Price Cap treatments.

In contrast, the bid quantities consistently exceeded optimal levels. Our findings reveal a substantial excess of demand at all valuation levels, resulting in an inefficient allocation of widgets. Consequently, participants' valuations did not have a significant influence on the final quantity of widgets they secured during the experiment.

Figure 8 presents the average bidding strategy in each treatment. Buyers at all valuation levels attempt to purchase roughly the same quantity in period 1. Table 2 shows that the decision to bid for one or two widgets is statistically independent of valuation in every

¹⁰A coding error in our oTree application prevented us from recording the period 1 bid prices for 4 of the 11 Free Market groups (32 participants across all 10 rounds, or 320 bid-price observations). The early version of the application overwrote the period 1 bid price with the period 2 bid price. This error does not affect transaction prices, quantities, or any other variable, which were recorded correctly for all groups and rounds. Bid-price figures and analyses for the Free Market treatment use only the 7 unaffected groups.

treatment: neither Pearson’s chi-squared test nor the Wilcoxon–Mann–Whitney test rejects equal distributions at the 5% level. In the Liberalization treatment, the Mann–Whitney test is marginally significant ($p = 0.060$), suggesting that high-value buyers may demand more, but the chi-squared test does not support this interpretation ($p = 0.166$). Bidding quantities therefore bear little relation to assigned valuations, contributing to excess demand and inefficient allocation.

[Figure 6 about here]

[Figure 7 about here]

[Figure 8 about here]

[Table 2 about here]

4.3 Market prices

Table 3 displays the average trading price in all treatments and trading periods. Prices fall within the expected range in the Free Market and Liberalization treatments. In the Anticipated Price Cap treatment, the first-period price (73.00) is above the myopic equilibrium but below the rationing equilibrium prediction of [88, 90].

In the Free Market treatment, prices in both periods align with the competitive equilibrium. The first-period price averages 73.67 (predicted: [70, 80]) and the second-period price averages 88.15 (predicted: [80, 90]).

In the Anticipated Price Cap treatment, the second-period price matches the binding price cap at 49.93. However, the first-period price is lower than predicted by the rationing equilibrium, suggesting that participants bid too little for storage.

In the Liberalization treatment, the first-period price matches the binding cap at 49.96. Once the market is liberalized in the second period, the price jumps to 109.04 — more than

double the first-period cap. This lies within the range $[90, 120]$ that carries 95% of the predicted price distribution, but the magnitude of the increase illustrates the consequences of the price-control-then-liberalize sequence. Second-period quantities drop to 2.75 (vs a predicted 3), as the free market clears at a much higher price with fewer participants.

Figure 9 illustrates the distribution of trading prices for each treatment and trading period over the last 5 rounds.¹¹ In most markets, prices consistently align with the equilibrium range. The most frequent price tends to be close to the optimal pricing range.

Table 5 compares the first-period price and storage across treatments. Participants in the Anticipated Price Cap treatment behave almost exactly like those in the Free Market treatment: neither first-period prices (73.67 versus 73.00, $p = 0.73$ across markets; across sessions, Mann–Whitney $p = 0.77$ and exact permutation $p = 0.88$) nor storage (2.36 versus 2.31 units per market, $p = 0.73$; across sessions, $p = 0.68$ and $p = 0.56$) differ. This null is not simply a consequence of low power. Even at the level of the six and nine independent sessions, the test can detect a 10.9-token difference in first-period prices at 80% power, less than the roughly 14-token gap predicted by the theory (rationing $[88, 90]$ versus competitive $[70, 80]$). With an anticipated cap, high-value buyers should bid aggressively ($p_1^* \in [88, 90]$) and store 3 units, rather than bid within $[70, 80]$ and store 2 units as in the Free Market treatment. Instead, Anticipated Price Cap buyers price storage as if the second-period market would remain competitive. The bidding quantities in Table 2 tell the same story: high-value buyers do not bid for more widgets than low-value buyers in any treatment. Participants respond to neither margin of the rationing equilibrium; they raise neither their prices nor their quantities. Storage in both treatments nevertheless exceeds that under liberalization (1.15 units per market), where storing requires obtaining two units in the rationed market, so realized storage is limited by the allocation lottery itself.

[Table 3 about here]

[Figure 9 about here]

¹¹Tables 11, 12, 13, and 14 in the appendix replicate the main results tables using all 10 rounds. The results are qualitatively identical.

Except in the Liberalization treatment, we find no prolonged learning or convergence toward equilibrium. Figures 18 and 17 show that average trading prices adjust between the first two rounds and change little thereafter. Table 7 likewise shows that quantities purchased are stable between the first and last five rounds. Only the Liberalization treatment exhibits a statistically significant movement toward equilibrium.

The treatments also differ in the feedback they provide. Under liberalization, the first-period cap produces an observable queue and a sharply higher second-period price; in the Free Market treatment, the fourfold cost increase is transparent. Under the Anticipated Price Cap, buyers must instead anticipate a future rationing probability that depends on the storage decisions of others. Feedback about this probability is delayed and noisy. This difference may explain why behavior converges in some treatments but not under the anticipated cap, although the experiment was not designed to test that mechanism.

[Table 5 about here]

4.4 Consumption and hoarding

Table 6 displays the average quantities traded in each market, separated by treatments and trading periods. In the Free Market treatment, the first-period quantity (7.04) is close to the predicted range of $[7, 8]$, while the second-period quantity (2.05) falls below the predicted 3 units, likely due to excessive storage. In the Liberalization treatment, traded quantities are close to the predicted values. In the Anticipated Price Cap treatment, the first-period quantity (7.09) is near the predicted 8 units, and the second-period quantity matches the prediction of 1 unit exactly.

As indicated previously, we do not observe a clear convergence path to the equilibrium. The results in Table 7 show that the amount of goods purchased is similar between the first and last five rounds for all treatments, except for the Liberalization treatment.

[Table 6 about here]

[Table 7 about here]

Traded quantities often differ from the competitive equilibrium. Figure 11 shows their distribution during the last five rounds. As the bidding results suggest, buyers pay attention to prices but appear to neglect the equilibrium consequences of the quantities they demand. Figure 12 also shows that valuation does not consistently determine demand: high-value buyers are sometimes excluded, while low-value buyers obtain goods in all treatments even when they have little to gain from doing so.

[Figure 11 about here]

[Figure 12 about here]

Figure 13 shows the number of goods stored in the first period. In the Free Market treatment, participants store slightly more than predicted (2.36 vs 2 units per market). In the Anticipated Price Cap treatment, participants store less than predicted (2.31 vs 3 units per market). The two treatments produce nearly identical storage levels despite different theoretical predictions, suggesting that participants respond similarly to storage incentives regardless of whether the motivation is competitive arbitrage or defensive hoarding. In the Liberalization treatment, participants store on average 1.15 units per market. Unlike the deterministic storage benchmarks in the other treatments, the liberalization benchmark is an expectation over the allocation lottery: a buyer stores only if both of her requested units succeed, so with six two-unit bidders expected storage is $6 \times (5/14)(4/13) \approx 0.66$ units per market. Observed storage is thus well above the full-demand benchmark rather than below it. The gap has a natural explanation: the benchmark assumes that all fourteen requested units compete in the lottery, whereas effective demand at the transaction price was thinner than predicted (the average shortage rate is 25%, against a predicted 64%), which raises the probability that a two-unit bid succeeds twice. Storage remains motivated by capturing the rent between the regulated first-period price and the free second-period price rather than by avoiding future scarcity. In contrast to the Anticipated Price Cap

treatment, then, participants do not under-store under liberalization: when the storage motive operates through a currently experienced shortage rather than an anticipated one, storage meets or exceeds the equilibrium benchmark. After rising over the first rounds, storage is broadly stable in the Free Market and Anticipated Price Cap treatments; in the Liberalization treatment it is noisier and, consistent with the worsening first-period shortage, drifts over the rounds rather than holding constant (Figure 13).

[Figure 13 about here]

4.5 Shortage and efficiency

Prices serve as signals that help people understand what is in demand and what is available. Efficiency is maximized when supply and demand determine prices. When external forces impose distortions on the price mechanism, these signals become distorted, misleading agents and causing inefficiencies. In this section, we analyze the impact on shortage and efficiency across our three treatments.

Storage can improve efficiency by insuring buyers against an anticipated distortion, as in the Anticipated Price Cap treatment. When a controlled market is expected to be liberalized, however, buyers have an incentive to purchase excess units at the low first-period price and capture rents. The Free Market treatment entails neither price controls nor predicted rationing.

Table 8 shows the average shortage rate for each treatment and period, using all 10 rounds.¹² We calculate the shortage rate as the share of unsatisfied demand among buyers who bid at or above the transaction price. The results confirm our main theoretical predictions. In the Free Market treatment, shortage rates are near zero in both periods (4% in period 1, 3% in period 2), consistent with a free market without binding price controls.

¹²We use all 10 rounds for the shortage analysis rather than the last 5 for two reasons. First, the worsening-over-rounds dynamic in the Liberalization treatment is itself a central finding that requires the full sample to document. Second, the coding error that affected period 1 bid prices in 4 Free Market groups reduces the available sample for that treatment; using all rounds preserves statistical power. As documented in Table 7, quantities are stable across early and late rounds, so this choice does not bias the shortage estimates.

In the Liberalization treatment, where the price cap binds in the first period, we observe a 25% average shortage rate in period 1. However, rates between 50% and 60% are the most frequently recorded shortage rate, as shown in Figure 14. The shortage worsens over rounds: the rate rises from 12% in round 1 to 34% in round 10. As discussed under H3a, the increase reflects more effective demand at the controlled price and is consistent with participants learning to position their bids to obtain regulated units before liberalization. This dynamic suggests that shortages under a binding price cap may intensify as market participants gain experience with the regulated environment. Once the market is liberalized in the second period, shortage drops to 2%, but at the cost of a sharp price increase.

In the Anticipated Price Cap treatment, the first period operates as a free market and shows a 9% shortage rate. The second period, where the price cap binds, shows a 60% shortage rate. Storage helps reduce the shortage below the static prediction, but does not eliminate it.

[Table 8 about here]

[Figure 14 about here]

To analyze efficiency, we measure total surplus: the sum of consumption values minus production costs and storage costs. Since sellers are automated, this captures the full welfare generated by the market. We define the efficiency ratio κ for market m as:

$$\kappa_m = \frac{W_m}{W_m^*}, \quad W_m = \sum_{i \in B} r_i q_{im} - \sum_{j \in S} c_j y_{jm} - \alpha_m s$$

where $r_i q_{im}$ is the consumption value of buyer i , $c_j y_{jm}$ is the production cost of seller j , α_m is the number of stored units, and $s = 10$ is the per-unit storage cost. The transaction price cancels out of total surplus because it is a transfer between buyers and sellers. The benchmark W_m^* is the theoretical equilibrium total surplus: the competitive equilibrium for the Free Market treatment and the rationing equilibrium for the two price-cap treatments.

Table 9 reports total surplus and the efficiency ratio. Total surplus is highest in the Free Market treatment (701 LCU per market), followed by the Anticipated Price Cap (647 LCU) and Liberalization (609 LCU). Efficiency ratios are 77%, 75%, and 84%, respectively. The Free Market and Anticipated Price Cap treatments each lose roughly a quarter of their theoretical surplus. The Liberalization ratio is higher, but against a much lower benchmark ($W^* = 915, 864, \text{ and } 723$): rationing and the storage lottery depress the attainable surplus itself, so participants capture a larger share of a smaller pie, and the absolute welfare losses under price controls remain large.

Realized total surplus is highest in the Free Market treatment and lowest under Liberalization, where rent-seeking storage predominates. In the Free Market treatment, participants store slightly more than predicted (2.36 rather than 2 units per market), while the second-period quantity falls below equilibrium (2.05 rather than 3). In the Anticipated Price Cap treatment, participants store less than the rationing equilibrium predicts (2.31 rather than 3) and bid too little for storage. Despite these deviations in opposite directions, storage is nearly identical in the two treatments. This similarity is not present at the outset: average storage in round 1 is 1.64 units in the Free Market treatment and 1.43 under the Anticipated Price Cap. It rises to 2.36 and 2.31 units, respectively, over rounds 6–10. At the individual level, participants store in 28.6% of rounds on average in the Free Market treatment and in 27.9% under the Anticipated Price Cap, and the distributions of individual storage frequencies are statistically indistinguishable (Mann–Whitney $p = 0.82$). Thus, although the additional rationing incentive in the Anticipated Price Cap treatment raises the theoretical benchmark from 2 to 3 units, it has no detectable effect on storage behavior.

The Liberalization treatment produces the lowest total surplus of all three treatments. This outcome results from the combination of three forces. First, the binding price cap in period 1 causes rationing, randomly excluding some buyers (including high-value ones) from the market. Second, buyers who manage to store at the regulated price capture large rents (paying 60 for goods worth 109 in the free market), while buyers rationed out of period 1 face the full liberalized price. Third, the second-period quantity collapses to 2.75 units (vs

a myopic prediction of 3), as the market clears at a much higher price.

These benchmarks isolate the price mechanism and do not incorporate risk aversion or preferences for early resolution of uncertainty; they therefore need not describe utility-maximizing behavior under such preferences.

[Table 9 about here]

Figure 15 compares the efficiency ratio across rounds for each treatment. The Free Market treatment shows the clearest learning trajectory, rising from 59% in round 1 to 85% by round 10. The Anticipated Price Cap treatment improves over the first few rounds and then fluctuates around 75% with no further gains, and the Liberalization treatment likewise improves early and then fluctuates around 84% over the last five rounds, suggesting that the efficiency losses from price controls persist with experience rather than diminish.

[Figure 15 about here]

4.6 Individual behavior and market power

Because the sellers are automated and play competitively, the only side of the market that can exercise market power is the buyers': the relevant concern is buyer-side monopsony, in which a buyer shades his bid to depress the clearing price. We examine this directly in the free first-period markets — the Free Market and Anticipated Price Cap treatments, whose first period is uncapped — by reconstructing the call-market clearing exactly from the submitted bids: we take the price to be the bid of the marginal accepted unit and apply the same clearing rules as the experimental software, including the token-budget constraint on bids. The reconstruction matches the observed clearing price in every market, so the individual bids let us measure each buyer's price impact precisely.¹³

¹³As with the bidding-strategy figures, the Free Market treatment uses only the four sessions (seven groups) with recorded first-period bids, while the Anticipated Price Cap treatment uses all nine sessions. Inference is at the level of the session, of which there are thirteen with usable bids; the individual-level regression below clusters on the session and reports wild-cluster-bootstrap p -values, since buyers interact repeatedly within a session. Clustering on the market group instead (21 clusters) gives the same estimates, with $p = 0.001$ and $p = 0.026$ for the two specifications reported below.

Buyers do shade their bids. The median markdown below valuation, $(r_i - b_i)/r_i$, is 0.26 in the Free Market treatment and 0.27 in the Anticipated Price Cap treatment — about a quarter of the assigned value — and it is positive in every one of the thirteen sessions (Wilcoxon $p < 0.001$). This shading lowers the clearing price by roughly 13 tokens relative to truthful bidding ($p < 0.001$). It does not by itself establish market power: in a uniform-price call market, even a buyer unable to move the price need not bid his full value. If buyers exercise market power, shading should increase when they are pivotal, that is, able to move the clearing price. A regression of the individual markdown on pivotality yields a strong positive coefficient (+0.162, wild-cluster-bootstrap $p = 0.001$), but this is a valuation artifact: high-value buyers are both more likely to be pivotal and shade more in proportional terms. Controlling for valuation and its square reverses the sign of the pivotality coefficient, to -0.076 ($p = 0.035$): conditional on valuation, pivotal buyers shade slightly less, not more, than buyers who cannot move the price. Buyers therefore do not exploit price-setting power; shading reflects valuation rather than the strategic exercise of monopsony.

Consistent with the absence of market power, the free-market clearing price stays at the competitive benchmark. In the Free Market treatment the first-period price is not significantly different from the midpoint of the competitive band (mean gap -1.6 tokens across the six sessions, Wilcoxon $p = 0.44$). This mirrors the strategic-wedge calculation of Section 2: whatever monopsony scope a thin single-period market might in principle offer is eliminated once storage thickens the market, and the data show no residual price manipulation.

5 Conclusions

This paper studies theoretically and experimentally how price distortions generate hoarding through intertemporal incentives. The motive for hoarding depends critically on the policy regime. When buyers anticipate a future price cap, they engage in defensive storage to insure against rationing. When they face a regulated market with anticipated liberalization, they engage in rent-seeking storage to exploit the gap between the regulated current price and the

higher future market price. Although both mechanisms shift purchases toward the present, their implications for prices, shortages, and welfare are quite different.

The broad patterns observed in the laboratory are consistent with the model. Shortages arise only when price caps bind; periods without controls exhibit near-zero shortage rates. In the Free Market treatment, participants store slightly more than the competitive prediction (2.36 versus 2 units). Under the anticipated cap, by contrast, storage falls short of the rationing-equilibrium benchmark (2.31 versus 3 units), while under liberalization realized storage (1.15 units) exceeds the expected 0.66 units implied by the allocation lottery, reflecting thinner-than-predicted effective demand in the rationed first period. The Free Market treatment generates the highest total surplus, while Liberalization, where rent-seeking rather than defensive storage predominates, generates the lowest.

The contrast between the Free Market and Anticipated Price Cap treatments is particularly informative. The introduction of a future price cap raises predicted storage from 2 to 3 units, yet produces no detectable change in either aggregate storage or the distribution of individual storage frequencies. Participants also bid substantially less for storage than the model predicts. The failure to respond to the anticipated cap therefore does not appear to reflect a general reluctance to store. Rather, participants respond insufficiently to the incentive created by future rationing.

The Liberalization treatment displays a different dynamic. Once the price ceiling is removed, the price more than doubles, from 50 to 109. At the same time, the shortage in the regulated first period grows from 12% in round 1 to 34% in round 10 as participants increasingly attempt to acquire units at the controlled price. Thus, experience with the market intensifies competition for goods whose regulated price lies below their anticipated future value. The resulting combination of current rationing, rent-seeking storage, and a sharp post-liberalization price increase makes liberalization the lowest-surplus environment in our experiment.

Our individual-level evidence also helps rule out strategic demand reduction as an explanation for the departures from equilibrium. Even in these small markets, with repeated

interaction and opportunities to store, we find no evidence that buyers systematically exploit their ability to affect the clearing price. Bid shading does not respond to a buyer’s ability to move the price once we control for valuation, and free-market prices remain at the competitive benchmark. Thus, the welfare losses generated by price controls in our experiment are imposed on a market that is already competitive rather than correcting buyer-side market power.

Risk preferences provide at most a partial explanation for the observed storage behavior. We elicited risk preferences using a Holt–Laury multiple price list after the market experiment. These data are noisy, with roughly 40% of participants exhibiting inconsistent switching patterns. Nevertheless, risk tolerance is modestly correlated with storage overall (Spearman $\rho = 0.16$, $p = 0.007$). The relationship is somewhat stronger in the Free Market ($\rho = 0.20$, $p = 0.056$) and Liberalization ($\rho = 0.22$, $p = 0.041$) treatments, where storage has a speculative component, than under the Anticipated Price Cap ($\rho = 0.09$, $p = 0.344$), where storage instead provides insurance against rationing, although the differences between treatments are not themselves statistically significant. Stable differences in measured risk tolerance therefore cannot account for the storage shortfall in the anticipated-cap treatment.

What accounts for that shortfall remains an open question. Buyers may recognize the direct benefit of storing without fully internalizing how aggregate storage affects future rationing. This interpretation is consistent with evidence that individuals underappreciate the equilibrium responses of others to policy changes (Dal Bó et al., 2018). Alternatively, behavior may reflect preferences over the timing of uncertainty resolution (Kreps and Porteus, 1978), inaccurate beliefs about rationing, limited attention to the future market, or other departures from standard decision making (Ellsberg, 1961; Kahneman et al., 1991; Alan and Ertac, 2015; Sunstein, 2000; de Clippel and Rozen, 2024). Because these deviations were not anticipated when we designed the experiment, our data can establish their empirical regularity but cannot discriminate among these mechanisms.

Several extensions follow naturally. Storage costs are fixed at $s = 10$ in our experiment, but the model predicts that higher storage costs weaken both defensive and rent-seeking

hoarding, while lower costs strengthen them. Future experiments could vary storage costs, elicit beliefs about rationing directly, or introduce strategic sellers to study supply withdrawal, quality adjustment, and joint learning by both sides of the market. More broadly, the mechanism studied here applies whenever current decisions allow agents to arbitrage distortions in intertemporal prices. This suggests applications beyond explicit price controls, including monetary environments in which restrictions on exchange or distorted relative prices generate incentives to shift transactions across time or markets, along the lines of Clower (1967) and Kemp and Long (1990). The broader lesson is that the consequences of price distortions depend not only on the transactions they prevent today, but also on the forward-looking behavior they induce in anticipation of tomorrow's market.

Table 1: Participants

	Free Market	Anticipated Price Cap	Liberalization
No. of participants	88	112	88
Gender			
Man	58%	57%	56%
Woman	40%	40%	42%
Non-binary	1%	1%	2%
NR	1%	2%	0%
Major			
STEM	58%	62%	59%
Non-STEM	26%	25%	19%
Economics	3%	1%	3%
NR	12%	12%	18%
Level of studies			
Graduate	53%	72%	76%
Undergraduate	44%	25%	20%
NR	2%	3%	3%

Table 2: Difference in player's valuation for those who bid for 1 or 2 widgets

Player's valuation	Bidding Quantity = 1	Bidding Quantity = 2	Total	Pearson's Chi-2 Test	Wilcoxon Mann Whitney Test
Free Market					
50	46	54	100		
70	29	24	53		
90	19	34	53	6.3251	-1.2290
110	26	27	53	Pr = 0.276	Pr = 0.210
140	41	67	108		
200	22	31	53		
Anticipated Price Cap					
50	57	74	131		
70	31	35	66		
90	31	36	67	10.9424	-1.7041
110	18	52	70	Pr = 0.053	Pr = 0.082
140	46	89	135		
200	27	41	68		
Liberalization					
50	48	57	105		
70	29	26	55		
90	24	29	53	7.8241	-1.8457
110	16	37	53	Pr = 0.166	Pr = 0.060
140	42	68	110		
200	20	35	55		

Note: The table uses the data for the last 5 rounds in each treatment

Table 3: Market Price

Treatment	Period	N	Price	se	95% CI	Equilibrium
Free Market	1	55	73.67	1.44	[70.85, 76.50]	[70, 80]
	2	55	88.15	1.59	[85.03, 91.26]	[80, 90]
Anticipated Price Cap	1	70	73.00	1.32	[70.42, 75.58]	[88, 90]
	2	70	49.93	0.07	[49.79, 50.07]	[50, 50]
Liberalization	1	55	49.96	0.04	[49.89, 50.03]	[50, 50]
	2	55	109.04	3.51	[102.16, 115.91]	[90, 120]

Note: Uses data from last 5 rounds

Table 4: No-Arbitrage Condition Test (Free Market)

$p_2 - p_1 - s$	
Mean gap	4.47
Std. Error	1.93
t -statistic	2.32
p -value	0.024
N	55

Note: One-sample t -test of $H_0: p_2 - p_1 - s = 0$, where $s = 10$ is the storage cost. Uses data from last 5 rounds of the Free Market treatment.

Table 5: Differences in market price and stored goods

	Free Market	Anticipated Price Cap	Liberalization	T-Test	Mann-Whitney
Num. Markets	55	70	55	-	-
Price in $t = 1$	73.67 (1.441)	73.00 (1.318)	49.96 (0.036)	0.3433 Pr = 0.7320	2030.5000 Pr = 0.5984
Stored Goods	2.36 (0.108)	2.31 (0.097)	1.15 (0.102)	0.3400 Pr = 0.7344	1987.0000 Pr = 0.7417

Note: Uses data from last 5 rounds. T-test and Mann-Whitney compare Free Market vs Anticipated Price Cap.

Table 6: Market Quantities

Treatment	Period	N	Q	se	95% CI	Equilibrium
Free Market	1	55	7.04	0.10	[6.84, 7.23]	[7, 8]
	2	55	2.05	0.11	[1.83, 2.28]	[3, 3]
Anticipated Price Cap	1	70	7.09	0.08	[6.93, 7.24]	[8, 8]
	2	70	1.00	0.00	[1.00, 1.00]	[1, 1]
Liberalization	1	55	4.93	0.04	[4.86, 5.00]	[5, 5]
	2	55	2.75	0.12	[2.52, 2.97]	[3, 3]

Note: Uses data from last 5 rounds

Table 7: Sold widgets distribution by treatment and rounds

Widgets in P1	Round 1-5	Round 6-10	Total	Chi-2 Test	Mann-Whitney
Free Market					
0	213	183	396	4.6129	-1.5061
1	105	127	232	Pr = 0.100	Pr = 0.106
2	122	130	252		
Anticipated Price Cap					
0	265	226	491	5.8590	-1.8041
1	145	172	317	Pr = 0.053	Pr = 0.053
2	150	162	312		
Liberalization					
0	262	232	494	6.1290	-1.4172
1	112	145	257	Pr = 0.047	Pr = 0.112
2	66	63	129		

Table 8: Shortage Rate by Period

Treatment	Period	N	Shortage Rate	se	95% CI
Free Market	1	70	0.04	0.008	[0.02, 0.05]
	2	70	0.03	0.012	[0.01, 0.05]
Anticipated Price Cap	1	140	0.09	0.009	[0.07, 0.11]
	2	140	0.60	0.019	[0.56, 0.64]
Liberalization	1	110	0.25	0.017	[0.21, 0.28]
	2	110	0.02	0.008	[0.00, 0.04]

Note: Uses all 10 rounds. Shortage rate is the share of unsatisfied demand among buyers who bid at or above the transaction price. Period 1 bid prices are unavailable for 4 Free Market groups from early sessions (the oTree application recorded only the period 2 bid price). Period 2 bids are unavailable for buyers who stored (bought 2 units in period 1).

Table 9: Total Surplus and Efficiency Ratio

Treatment	N	Total Surplus	W^*	Efficiency (κ)	95% CI
Free Market	55	700.7	915.0	0.77	[0.73, 0.80]
Anticipated Price Cap	70	646.8	864.0	0.75	[0.72, 0.78]
Liberalization	55	609.2	722.8	0.84	[0.80, 0.89]

Note: Total surplus $W = \sum_i r_i q_i - \sum_j c_j y_j - \alpha s$, where $r_i q_i$ is the consumption value of buyer i , $c_j y_j$ is the production cost of seller j , α is the number of stored units, and $s = 10$ is the storage cost. W^* is the theoretical equilibrium total surplus. Uses data from last 5 rounds.

Free Market




Figure 3: Free Market. The figure displays the supply and demand schedules for both periods. In period 1, sellers' costs are $c_1(i)$ and the market operates freely. In period 2, costs increase fourfold to $c_2(i) = 4 \times c_1(i)$. The dashed horizontal lines indicate the competitive equilibrium price range.

Anticipated Price Cap




Figure 4: Anticipated Price Cap. The figure displays the supply and demand schedules for both periods. In period 1, sellers' costs are $c_1(i)$ and the market operates freely; the dashed horizontal lines indicate the rationing equilibrium price range $p_1^* \in [88, 90]$. In period 2, costs increase fourfold to $c_2(i) = 4 \times c_1(i)$ and a price ceiling $\bar{p}_2 = 50$ is imposed (dash-dotted line). Only one seller has costs below the ceiling in period 2.

Liberalization




Figure 5: Liberalization. The figure displays the supply and demand schedules for both periods. In period 1, sellers' costs are $c_1(i)$ and a price ceiling $\bar{p}_1 = 50$ is imposed (dash-dotted line), creating a shortage. In period 2, costs increase fourfold to $c_2(i) = 4 \times c_1(i)$ and the market is liberalized; the dashed horizontal lines indicate the range $[90, 120]$ on which the predicted price distribution concentrates.




Figure 6: Bid prices in period 1. The figure presents the average bid price in each treatment. The lines represent 90% confidence intervals. The dashed horizontal lines denote the equilibrium price ranges.




Figure 7: Bid prices in period 2. The figure presents the average bid price in each treatment. The vertical lines represent 90% confidence intervals. The dashed horizontal lines denote the equilibrium price ranges.




Figure 8: Bid quantities in period 1. The figure presents the average quantity bid in period 1 for each treatment. The vertical lines represent 90% confidence intervals. The dashed horizontal line at 1 indicates that the participant does not attempt to store.




Figure 9: Trading prices, last 5 rounds. The dashed vertical lines represent the competitive equilibrium for the Free Market treatment. For the two price-cap treatments, they indicate the rationing equilibrium.




Figure 10: Imperfect arbitrage in the free market. Distribution of the per-market arbitrage gap $p_2 - p_1 - s$ in the Free Market treatment (last 5 rounds, $N = 55$). Under no-arbitrage the gap is zero (solid line); the distribution lies to the right of zero, with a positive mean (dashed line), indicating imperfect arbitrage.




Figure 11: Traded quantities, last 5 rounds. The dashed vertical lines represent the competitive equilibrium for the Free Market treatment. For the two price-cap treatments, they indicate the rationing equilibrium.




Figure 12: Acquired widgets in period 1 by player's valuation and treatment. The figure presents the distribution of players based on the number of widgets acquired in period 1 across all rounds, categorized according to their valuation.




Figure 13: Storage market in all treatments. The figure shows the number of widgets stored in each round for the three treatments in the experiment. The dashed line represents the theoretical prediction for each market.




Figure 14: Shortage rate in the Liberalization treatment (period 1)




Figure 15: Total surplus efficiency ratio over rounds. The graphs represent the ratio of realized total surplus to theoretical total surplus for each treatment across rounds. The dashed line at 1.0 marks the theoretical benchmark.

A Additional tables and figures

A.1 Learning and valuation effects

Table 10: Widget's value distribution by treatment, bidding quantity and rounds

Player's valuation	Round 1-5	Round 6-10	Total	Pearson's Chi-2 Test
Free Market				
Bidding Quantity in P1 = 0				
50	5	10	15	6.0639 Pr = 0.300
70	3	2	5	
90	0	2	2	
110	4	2	6	
140	5	2	7	
200	1	2	3	
Bidding Quantity in P1 = 1				
50	47	46	93	1.9549 Pr = 0.855
70	23	29	52	
90	22	19	41	
110	22	26	48	
140	48	41	89	
200	25	22	47	
Bidding Quantity in P1 = 2				
50	58	54	112	1.5656 Pr = 0.905
70	29	24	53	
90	33	34	67	
110	29	27	56	
140	57	67	124	
200	29	31	60	

Table 10: Widget's value distribution by treatment, bidding quantity and rounds
(continued)

Player's valuation	Round 1-5	Round 6-10	Total	Pearson's Chi-2 Test
Anticipated Price Cap				
Bidding Quantity in P1 = 0				
50	8	9	17	3.0965 Pr = 0.542
70	3	4	7	
90	8	3	11	
110	0	0	0	
140	3	5	8	
200	3	2	5	
Bidding Quantity in P1 = 1				
50	63	57	120	5.1788 Pr = 0.394
70	26	31	57	
90	30	31	61	
110	32	18	50	
140	59	46	105	
200	24	27	51	
Bidding Quantity in P1 = 2				
50	69	74	143	2.7621 Pr = 0.737
70	41	35	76	
90	32	36	68	
110	38	52	90	
140	78	89	167	
200	43	41	84	

Table 10: Widget's value distribution by treatment, bidding quantity and rounds
(continued)

Player's valuation	Round 1-5	Round 6-10	Total	Pearson's Chi-2 Test
Liberalization				
Bidding Quantity in P1 = 0				
50	4	5	9	
70	1	0	1	
90	1	2	3	2.5990
110	3	2	5	Pr = 0.627
140	0	0	0	
200	1	0	1	
Bidding Quantity in P1 = 1				
50	54	48	102	
70	27	29	56	
90	26	24	50	2.8923
110	28	16	44	Pr = 0.717
140	56	42	98	
200	23	20	43	
Bidding Quantity in P1 = 2				
50	52	57	109	
70	27	26	53	
90	28	29	57	2.1286
110	24	37	61	Pr = 0.831
140	54	68	122	
200	31	35	66	

Notes: The table above provides an overview of the distribution of the value of the widgets during the first and last five rounds of the period for each of the evaluated treatments and for each quantity intended to acquire in period 1.

A.2 Robustness: All 10 Rounds

The main-text tables report results from the last 5 rounds (rounds 6–10) to allow for learning. The tables below replicate the key results using all 10 rounds. The qualitative conclusions are unchanged: prices stay at or near their predicted ranges (the Free Market second-period average exceeds its range by only 0.3 tokens, and the Anticipated Price Cap first period remains below the rationing prediction, as in the main text), the shortage pattern is identical, storage levels move by less than one quarter of a unit, and the efficiency ranking across treatments is preserved. No conclusion in the main text is reversed by using all ten rounds.

Table 11: Market Price

Treatment	Period	N	Price	se	95% CI	Equilibrium
Free Market	1	110	71.44	1.24	[69.01, 73.86]	[70, 80]
	2	110	90.28	1.99	[86.38, 94.18]	[80, 90]
Anticipated Price Cap	1	140	71.04	0.98	[69.12, 72.95]	[88, 90]
	2	140	49.82	0.11	[49.61, 50.03]	[50, 50]
Liberalization	1	110	48.51	0.37	[47.79, 49.23]	[50, 50]
	2	110	106.72	2.75	[101.33, 112.10]	[90, 120]

Note: Uses data from all 10 rounds

Table 12: Market Quantities

Treatment	Period	N	Q	se	95% CI	Equilibrium
Free Market	1	110	6.69	0.10	[6.50, 6.89]	[7, 8]
	2	110	1.92	0.08	[1.76, 2.08]	[3, 3]
Anticipated Price Cap	1	140	6.72	0.09	[6.55, 6.89]	[8, 8]
	2	140	1.00	0.00	[1.00, 1.00]	[1, 1]
Liberalization	1	110	4.68	0.05	[4.59, 4.77]	[5, 5]
	2	110	2.37	0.09	[2.20, 2.55]	[3, 3]

Note: Uses data from all 10 rounds

Table 13: Differences in market price and stored goods

	Free Market	Anticipated Price Cap	Liberalization	T-Test	Mann-Whitney
Num. Markets	110	140	110	-	-
Price in $t = 1$	71.44 (1.235)	71.04 (0.979)	48.51 (0.367)	0.2576 Pr = 0.7969	7908.0000 Pr = 0.7127
Stored Goods	2.29 (0.078)	2.23 (0.076)	1.17 (0.067)	0.5659 Pr = 0.5720	8014.5000 Pr = 0.5560

Note: Uses data from all 10 rounds. T-test and Mann-Whitney compare Free Market vs Anticipated Price Cap.

Table 14: Total Surplus and Efficiency Ratio

Treatment	N	Total Surplus	W^*	Efficiency (κ)	95% CI
Free Market	110	673.6	915.0	0.74	[0.71, 0.76]
Anticipated Price Cap	140	615.9	864.0	0.71	[0.69, 0.74]
Liberalization	110	565.2	722.8	0.78	[0.75, 0.82]

Note: Total surplus $W = \sum_i r_i q_i - \sum_j c_j y_j - \alpha s$, where $r_i q_i$ is the consumption value of buyer i , $c_j y_j$ is the production cost of seller j , α is the number of stored units, and $s = 10$ is the storage cost. W^* is the theoretical equilibrium total surplus. Uses data from all 10 rounds.

A.3 Robustness: Session-level inference

The hypothesis tests in the main text use the session as the unit of observation (six sessions in the Free Market treatment, nine in the Anticipated Price Cap treatment, and seven under Liberalization). Each session-round contains one market of eight buyers in 8-participant sessions and two in 16-participant sessions, giving eleven, fourteen, and eleven market groups in the three treatments (55, 70, and 55 markets over rounds 6–10). In the 16-participant sessions, however, the two groups of eight are re-formed at random each round, so the two market groups in such a session are neither fixed sets of buyers nor independent of each other. Table 15 reports, for each test in the text, the session-level Wilcoxon or Mann–Whitney p -value, an exact permutation p -value on the same session means, and the corresponding test that treats the market groups as independent observations. The permutation tests use the magnitudes rather than the ranks of the session means: for comparisons between treatments they enumerate every reassignment of the fifteen sessions to the two treatments (the randomization distribution implied by session-level treatment assignment), and for comparisons with a theoretical benchmark they enumerate every sign flip of the session deviations from the benchmark. Every conclusion is the same under either unit and under either test. The session-level tests are naturally less powerful: with six observations the smallest attainable two-sided Wilcoxon p -value is 0.031, and two results that are borderline at the group level (storage against the competitive level of 2, and the Free Market arbitrage gap) are not significant at the session level, as the main text notes.

Table 15: Main hypothesis tests: session level (main text) versus market-group level

Test	Treatment(s)	N sessions	Session rank test p	Session permutation p	Group-level p
p_1 vs. ceiling of 50 (H2b)	Anticipated Price Cap	9	0.004	0.004	0.001
Storage vs. ra- tioning bench- mark 3 (H2d)	Anticipated Price Cap	9	0.004	0.004	0.002
Storage vs. com- peti- tive bench- mark 2 (H2d)	Anticipated Price Cap	9	0.109	0.094	0.040
p_1 : FM vs. APC	FM, APC	6, 9	0.77	0.88	0.73
Storage FM vs. APC	FM, APC	6, 9	0.68	0.56	0.73
Arbitrage gap $p_2 -$ $p_1 - s$ vs. 0 (H1)	Free Market	6	0.219	0.156	0.083
p_2 vs. my- opic 110, one- sided (H3c)	Liberalization	7	0.766	0.727	0.416
Period- 1 short- age	Liberalization	7	0.016	0.016	0.001

A.4 Supplementary figures




Figure 16: No more than 6 cans of tuna per person, thank you

Note: Picture taken from Al Jazeera (2014), retrieved on October 10, 2023.

Figures 17–19 display the evolution of market prices and quantities over rounds for each treatment. In each figure, the first panel shows the market-clearing price per period and the second panel shows the quantities traded. Dashed horizontal lines represent theoretical predictions; when two lines are present, they indicate the upper and lower limits of the equilibrium range.




Figure 17: Free Market treatment outcomes




Figure 18: Anticipated Price Cap treatment outcomes




Figure 19: Liberalization treatment outcomes

B Experimental procedures

B.1 Description

The experiment is conducted in person. Once subjects arrive at the lab in Enterprise Hall, they are given a printed copy of the “Informed Consent Form”. If they verbally agree to the terms, they are seated in individual spots and are ready to begin the session. If they do not agree, they receive the show-up fee and leave.

The experiment is implemented using Otree 5.1, a web-based experimental software. Participants find the instructions and tasks on their computer screens. All answers they provide are anonymous, and participants cannot exchange any information with one another.

The experiment consists of 10 rounds. At the beginning of each round, subjects are

randomly divided into groups of eight participants, each corresponding to a market. All members are assigned the buyer role and are tasked with buying intangible widgets from nine automated sellers. They all receive the same budget of tokens to pay for the widgets.

Every round consists of two periods. In the first period, participants choose whether to buy zero, one, or two widgets and offer a price for them. Buying widgets lets them accumulate points, based on a randomly assigned valuation parameter. If a buyer secures two widgets, she receives points for both the first and second periods but must pay a price equal to 10 tokens, representing a “storage fee” for the second widget. If a buyer secures one widget, she receives only points for the first period and does not incur a storage fee. In the second period, those who didn’t buy enough widgets in the first period can purchase an additional widget without incurring a storage fee.

Trading employs a “single-call market” mechanism, a type of auction that resembles how electricity markets function in most states. In a single-call market, all market participants submit their bids and offers before the auction. The auctioneer then conducts a single round of bidding to match buyers and sellers at the clearing price. The clearing price is determined when the quantity of the good being auctioned matches the total quantity demanded and supplied.

Not all bids can be satisfied; it depends on the offered prices and the quantities all group members demand. In each period, the market can offer a maximum of 9 widgets. If demand exceeds supply, widgets are allocated randomly with an equal probability among those who made acceptable offers.

At the end of the second period, every participant in the group is informed about how many widgets they were able to purchase, the transaction price, and their earnings. The final payoff is determined by a randomly selected round, combined with the show-up fee and a bonus question.

There are three treatments, which vary based on the maximum possible price participants can offer and the parameters set for the automated sellers. Table 16 summarizes the treatment parameters.

Table 16: Treatment parameters

Treatment	Period-1 costs	Period-2 costs	Price restriction
Free Market	$\mathbf{c}_1$	$4\mathbf{c}_1$	$\max\{P_1, P_2\} = 200$
Anticipated Price Cap	$\mathbf{c}_1$	$4\mathbf{c}_1$	$P_2 \leq 50$
Liberalization	$\mathbf{c}_1$	$4\mathbf{c}_1$	$P_1 \leq 50$

Note: Period-1 seller costs are $\mathbf{c}_1 = (5, 20, 20, 30, 50, 55, 60, 80, 100)$ in all treatments.

Once the experiment finishes, there is a bonus question to assess the participants’ risk aversion. This question uses a multiple price list design. The subject is presented with 9 cases. For each case, the subject decides between one of two presented lotteries to participate in. These lotteries offer points with varying probabilities. At the end of this section, the

software randomly selects one case, and the subject participates in the chosen lottery, winning the corresponding extra payoff.

Each session lasts less than 60 minutes, including payments. In addition to the \$10 show-up fee, subjects can earn between \$8 and \$15. These calculations factor in the bonus question. Each subject is paid privately at a rate of 40 points = 1 USD.

B.2 Participation

The Interdisciplinary Center for Economic Science (ICES) maintains a pre-registered participant pool through our recruitment system at experiments.gmu.edu. We require all participants to be at least 18 years old and to be members of the Mason community (e.g., students, staff). Any adult aged 18 and over, regardless of gender, who is pre-registered in our system, is eligible to participate in the study. We do not impose any additional exclusions.

- Duration per session: 60 minutes
- Participants per session: 8 or 16

C Experimental Instructions

Instructions for the experiment are provided via video. The scripts below reproduce the content of each video. Because the three treatments share a common structure, we first present the full instructions for the Free Market treatment (Section C.1) and then describe only the treatment-specific modifications for the Anticipated Price Cap (Section C.2) and liberalization (Section C.3) treatments.

C.1 Free Market treatment

Slide 1: Welcome! Hello! Welcome to our economics study.

Slide 2: Overview This video explains the two sections in this session: a market game and a bonus question with a short survey. We'll focus mainly on the market game.

Slide 3: Market Game Introduction You will play a market game with two trading periods: Period one and Period two.

- You will act as a buyer.
- Your job is to buy imaginary objects called widgets at the best price to maximize your earnings.

Slide 4: Market Structure You will join seven other buyers and nine computer sellers to form the market.

- Each seller can only offer one widget in each trading period.
- At most nine widgets are available per period.

Slide 5: Utility Points and Rounds

- You will be assigned a random value between 50 and 200 utility points.
- This value is the reward you get when you buy a widget, and it changes each round.
- There will be 10 rounds in total.
- At the end, utility points will be converted to cash at a rate of 40 points to 1 dollar.

Slide 6: Bidding Process Here's how the market works:

- You have a budget of 200 tokens to bid for widgets.

Slide 7: Market Mechanics

- Buyers submit bids for widgets, and sellers submit asking prices.
- The auctioneer will find the “market-clearing price” - the lowest price where all available widgets are sold.

- Your bid is the token amount you're willing to pay for a widget.
- If a seller's asking price is higher than the market-clearing price, they cannot produce the widget.
- All sellers will keep their asking prices constant.
- **In this experiment, sellers will experience an increment in production costs and their asking prices will be four times higher in the second period.**

Slide 8: Widget Purchase All buyers who bid at or above the market-clearing price will successfully buy a widget.

Slide 9: Example In the example you see on the screen, seven sellers produced the widget and only seven buyers made acceptable offers. Note that there are no unsatisfied acceptable offers.

Slide 10: Demand and Supply

- There can be situations in an auction when the demand for widgets is greater than the supply.
- In such cases, widgets will be randomly allocated among all buyers who made acceptable offers.

Slide 11: Earnings and Costs

- If you successfully buy a widget, you will earn utility points from the widget of that period.
- The market-clearing price will be deducted from your budget.
- If your bid was not accepted, you won't pay any price.

Slide 12: Random Allocation

- In cases that demand exceeds the supply, the offer will be satisfied randomly.
- In our previous example, there are 2 demanded widgets for each available widget, implying that each acceptable offer will be satisfied with a probability of $1/2$.

Slide 13: Trading and Storage

- In the first period, you can buy either zero, one or two widgets.
- If you buy two widgets, you will need to pay a storage fee of 10 tokens, and one widget will be saved in your account.
- If you buy one or no widgets in the first period, you will need to bid again in the second period.

Slide 14: Strategies and Rewards

- To help you maximize your earnings, remember to bid strategically and observe the market dynamics.
- We will repeat the market game for ten rounds.
- At the end of each round, your rewards and any unused tokens will be converted to US dollars at a rate of 40 tokens to 1 dollar.

Slide 15: Second Section The second section of the experiment is a bonus question where you will choose between two lotteries for nine different cases. One of the cases will be randomly selected, and the lottery you chose will be added to your earnings.

Your final payment will be the sum of the show-up fee, game reward, and bonus question reward.

If you have any questions or need clarification, please don't hesitate to ask the experimenter.

Good luck with the experiment and thank you for your participation!

C.2 Anticipated Price Cap treatment

The Anticipated Price Cap treatment uses the same instructions as the Free Market treatment (Section C.1), with the following modification:

Slide 7 (Market Mechanics) adds the following bullet after the announcement about seller costs increasing fourfold:

- **To allow players with lower valuations to have access to widgets, a maximum price of 50 tokens is imposed in the second period.**

All other slides are identical.

C.3 Liberalization treatment

The Liberalization treatment uses the same instructions as the Free Market treatment (Section C.1), with the following modification:

Slide 7 (Market Mechanics) replaces the bullet about seller costs increasing fourfold with:

- **In this experiment, sellers will experience an increment in production costs and their asking prices will be four times higher in the second period. To allow players with lower valuations to have access to widgets, a maximum price of 50 tokens is imposed in the first period. In the second period, there is no price restriction.**

All other slides are identical.

D Comprehension Quiz

After the instructions, the following quiz was delivered in a timed page within the online application. Participants had 15 minutes to answer all questions. They were allowed to ask the experimenter questions while completing it, so the quiz served to ensure comprehension rather than to screen participants; no participant was excluded on the basis of quiz performance. Performance was nonetheless high: across the 288 participants, the average score was 84% of the fourteen items correct, the median participant answered every item correctly, and 84% scored at least 80%. This indicates that participants understood the market environment before trading began.

1. What is the benefit of buying widgets?
 - a) You obtain additional tokens.
 - b) You obtain utility units that you can convert to cash at the end.
 - c) None of the above.
2. If you buy no widgets, what will happen?
 - a) I will have no rewards.
 - b) I will get a widget randomly.
 - c) The tokens that I did not use will get converted into cash at the end.
3. Do the benefits that buyers get for buying widgets stay the same during the two periods in each round?
 - a) False
 - b) True
4. Do buyers' rewards from buying widgets remain the same for all rounds?
 - a) False
 - b) True
5. Do sellers' production costs remain the same for both periods within the round?
 - a) False
 - b) True
6. At what rate will utility units be converted into cash?
 - a) 60 units to 1 USD
 - b) 30 units to 1 USD
 - c) 40 units to 1 USD.

7. At what rate will tokens be converted into cash?
- a) 60 units to 1 USD
 - b) 30 units to 1 USD
 - c) 40 units to 1 USD.
8. How many widgets are the maximum number available in the market?
- a) 8
 - b) 9
 - c) 10.
9. Are there two periods in which you can buy an extra widget to store in each period?
- a) False
 - b) True.
10. What is the cost of using storage for a widget?
- a) 0 tokens.
 - b) 10 tokens.
 - c) 20 utility units.
11. What needs to happen to have a successful purchase?
- a) The price offered must be greater than or equal to the market price.
 - b) The available widgets in the market are less than the total number of widgets demanded at the market price.
 - c) Both of the above.
12. If there are bids for 6 widgets at a price of 100, but only 3 sellers are willing to sell at that price, the 3 widgets from the sellers will be allocated randomly. With what probability will a bid be successful?
- a) $1/2$
 - b) 1
 - c) $1/6$.
13. Will you be able to offer any price within your budget?
- a) Yes.
 - b) No, that is only in the first period. In the second period, there is a price cap of 50 tokens.

- c) No, you can offer any amount regardless of the budget.
14. Which round will determine your final payoff?
- a) The first one - Round 1.
 - b) The last one - Round 10.
 - c) It will be selected randomly by the computer at the end of the experiment.

Declaration of generative AI and AI-assisted technologies in the manuscript preparation process

During the preparation of this work the authors used Claude (Anthropic) and ChatGPT 5.6 in order to assist with copy editing, text organization, and verification of references. After using these tools, the authors reviewed and edited the content as needed and take full responsibility for the content of the published article.

References

- Al Jazeera (2014). Is hoarding causing Venezuela food shortages? Available from: <https://www.aljazeera.com/features/2014/3/2/is-hoarding-causing-venezuela-food-shortages>, Accessed: October 9, 2023.
- Alan, S. and Ertac, S. (2015). Patience, self-control and the demand for commitment: Evidence from a large-scale field experiment. Journal of Economic Behavior & Organization, 115:111–122. Behavioral Economics of Education.
- Archivo DiFilm (1991). Informe precios de los combustibles: Aumento de la Nafta-Gasol 1990. Buenos Aires, Argentina. Available from: <http://www.difilm-argentina.com>, Accessed: October 23, 2023.
- Arrow, K. J., Harris, T., and Marschak, J. (1951). Optimal inventory policy. Econometrica, 19(3):250–272.
- Brekke, K. R., Nuscheler, R., and Rune Straume, O. (2006). Quality and location choices under price regulation. Journal of Economics & Management Strategy, 15(1):207–227.
- Cason, T. N. and Friedman, D. (1997). Price formation in single call markets. Econometrica, 65(2):311–345.
- Chen, D. L., Schonger, M., and Wickens, C. (2016). oTree—an open-source platform for laboratory, online, and field experiments. Journal of Behavioral and Experimental Finance, 9:88–97.
- Clower, R. (1967). A reconsideration of the microfoundations of monetary theory. Economic Inquiry, 6(1):1–8.
- Coursey, D. L. and Smith, V. L. (1983). Price controls in a posted offer market. The American Economic Review, 73(1):218–221.
- Coyne, C. and Coyne, R. (2015). Flaws and ceilings : price controls and the damage they cause. Institute of Economic Affairs, London, UK.

- Crawford, V. P., Costa-Gomes, M. A., and Iriberri, N. (2013). Structural models of nonequilibrium strategic thinking: Theory, evidence, and applications. Journal of Economic Literature, 51(1):5–62.
- Dal Bó, E., Dal Bó, P., and Eyster, E. (2018). The demand for bad policy when voters underappreciate equilibrium effects. Review of Economic Studies, 85(2):964–998.
- de Clippel, G. and Rozen, K. (2024). Bounded rationality in choice theory: A survey. Journal of Economic Literature, 62(3):995–1039.
- Deacon, R. T. and Sonstelie, J. (1989). The welfare costs of rationing by waiting. Economic Inquiry, 27(2):179–196.
- Department of Justice (2022). Coronavirus response: Combating price gouging & hoarding. Accessed: October 9, 2023.
- Drèze, J. H. (1975). Existence of an exchange equilibrium under price rigidities. International Economic Review, 16(2):301–320.
- Ellsberg, D. (1961). Risk, ambiguity, and the Savage axioms. Quarterly Journal of Economics, 75(4):643–669.
- Finley, G., Holt, C., and Snow, E. (2019). The welfare costs of price controls and rent seeking in a class experiment. Experimental Economics, 22(3):753–771.
- Guenette, J. D. (2020). Price controls: Good intentions, bad outcomes. Policy Research Working Paper 9212, World Bank.
- Hansman, C., Hong, H., de Paula, A., and Singh, V. (2020). A sticky-price view of hoarding. NBER Working Paper 27051, National Bureau of Economic Research. Revised September 2021.
- Hatfield, J. W., Plott, C. R., and Tanaka, T. (2012). Understanding price controls and nonprice competition with matching theory. American Economic Review, 102(3):371–75.
- Hatfield, J. W., Plott, C. R., and Tanaka, T. (2016). Price controls, non-price quality competition, and the nonexistence of competitive equilibrium. Games and Economic Behavior, 99:134–163.
- Hayek, F. A. (1945). The use of knowledge in society. American Economic Review, 35(4):519–530.
- Herings, P. J.-J. (2018). Equilibrium and matching under price controls. Journal of Economic Theory, 177:222–244.
- Isaac, R. M. and Plott, C. R. (1981). Price controls and the behavior of auction markets: An experimental examination. American Economic Review, 71(3):448–459.

- Ivanova, Y. and Sarmiento, A. (2013). Evaluaci'ón de la huella h'ídrica de la ciudad de Bogotá como una herramienta de la gesti'ón del agua en el territorio urbano. Revista Ambiental Agua, Aire y Suelo, 4(2):1–5.
- Kahneman, D., Knetsch, J. L., and Thaler, R. H. (1991). Anomalies: The Endowment Effect, Loss Aversion, and Status Quo Bias. Journal of Economic Perspectives, 5(1):193–206.
- Kemp, M. C. and Long, N. V. (1990). The relaxation of price controls, the general price level, and welfare. Journal of Macroeconomics, 12(2):221–246.
- Kreps, D. M. and Porteus, E. L. (1978). Temporal resolution of uncertainty and dynamic choice theory. Econometrica, 46(1):185–200.
- Madden, P. J. (1984). Existence of Drèze equilibrium under set-up costs. Journal of Economic Theory, 33(2):275–288.
- Martinelli, C., Wang, J., and Zheng, W. (2023). Competition with indivisibilities and few traders. Experimental Economics, 26(1):78–106.
- National Public Radio (2023). Why people in Argentina are in a race against time to spend their earnings. Accessed: October 9, 2023.
- Nguyen, T. T. and Whalley, J. (1986). Equilibrium under price controls with endogenous transactions costs. Journal of Economic Theory, 39(2):290–300.
- Rydzewski, M. (2020). Price gouging during a pandemic: The federal government's response. DTTP: Documents to the People, 48(3):10–15.
- Smith, V. L. and Williams, A. W. (1981). On Nonbinding Price Controls in a Competitive Market. American Economic Review, 71(3):467–474.
- Sunstein, C. R. (2000). Cognition and cost-benefit analysis. Journal of Legal Studies, 29(S2):1059–1103.
- The White House (2025). Fact sheet: President Donald J. Trump addresses security risks from price fixing and anti-competitive behavior in the food supply chain. Accessed: March 2026.
- Vallejo, M., editor (2011). Talleres de crónica: Memorias del agua en Bogotá. Antología. Alcaldía Mayor de Bogotá, Imprenta Distrital, Bogotá.
- Wall Street Journal (2016). Venezuelans, facing food shortages, rally behind vilified conglomerate. The Wall Street Journal, June 3, 2016. Available from: https://www.wsj.com/articles/venezuelas-biggest-private-company-fights-for-survival-1464964360?mod=fox_australian, Accessed: October 9, 2019.

- Wedig, G., Mitchell, J. B., and Cromwell, J. (1989). Can price controls induce optimal physician behavior? Journal of Health Politics, Policy and Law, 14(3):601–620.
- Weitzman, M. L. (1974). Prices vs. quantities. The Review of Economic Studies, 41(4):477–491.
- Weitzman, M. L. (1991). Price distortion and shortage deformation, or what happened to the soap? American Economic Review, 81(3):401–414.
- Wilkie, C. (2024). Harris to propose federal ban on ‘corporate price-gouging’ in food and groceries. CNBC. Accessed: March 2026.
- Yohe, G. W. (1978). Towards a general comparison of price controls and quantity controls under uncertainty. The Review of Economic Studies, 45(2):229–238.