Chaos and Misallocation under Price Controls

Brian C. Albrecht[*], Alex Tabarrok[†], and Mark Whitmeyer[‡]

February 11, 2026

Price controls kill the incentive for arbitrage. We prove a Chaos Theorem: under a binding price ceiling, suppliers are indifferent across destinations, so arbitrarily small cost differences can determine the entire allocation. The economy tips to corner outcomes in which some markets are fully served while others are starved; small parameter changes flip the identity of the corners, generating discontinuous welfare jumps. These corner allocations create a distinct source of cross-market misallocation, separate from the aggregate quantity loss (the Harberger triangle) and from within-market misallocation emphasized in prior work. They also create an identification problem: welfare depends on demand far from the observed equilibrium. We derive sharp bounds on misallocation that require no parametric assumptions. In an efficient allocation, shadow prices are equalized across markets; combined with the adding-up constraint, this collapses the infinite-dimensional welfare problem to a one-dimensional search over a common shadow price, with extremal losses achieved by piecewise-linear demand schedules. Calibrating the bounds to station-level AAA survey data from the 1973–74 U.S. gasoline crisis, misallocation losses range from roughly 1 to 9 times the Harberger triangle.

JEL Classification: D45, D61, L51, Q41

Keywords: Price controls; Misallocation; Deadweight loss; Market segmentation; Shadow prices; Gasoline shortages

[*]International Center for Law & Economics. Email: mail@briancalbrecht.com.
[†]George Mason University. Email: Tabarrok@gmu.edu.
[‡]Arizona State University. Email: mark.whitmeyer@gmail.com.

1 Introduction

In February 1974, the aggregate gasoline shortfall in the United States was roughly 9 percent nationwide.[1] Price controls on oil kept the price at the pump below the market price. If the standard model of price ceilings held and markets are similar, each local market would have experienced a roughly proportional 9 percent reduction, shadow prices would equalize, and the familiar Harberger triangle would measure the welfare cost. Instead, pain fell in a patchwork. As Yergin (1991) describes, ”gasoline was in short supply in major urban areas, but there were more than abundant supplies in rural and vacation areas, where the only shortage was of tourists.” Why should a 9 percent national shortfall leave some markets with nothing while others have plenty? This paper shows that feast or famine, not modest belt-tightening, is the generic outcome of price controls.

The argument rests on a simple observation: price controls kill the incentive for arbitrage. Under market prices, if gasoline is more valuable in Connecticut than in Idaho, traders profit by redirecting supply until shadow prices equalize. A slight cost advantage for one destination captures only the marginal tanker loads, because the price premium in the underserved market rises just enough to offset the cost difference. But when prices are frozen at $\bar{p}$ everywhere, there is no price gap to exploit. Every unit earns identical revenue regardless of destination, so suppliers are indifferent about where to ship. A one-cent cost advantage is as good as a one-dollar advantage: both redirect the entire flow, because there is no price in the destination market pushing back. What would have been a marginal reallocation becomes a categorical one, and the economy loses its capacity for incremental adjustment, instead lurching between all-or-nothing extremes (Sowell 1980, p. 137). Allocation is no longer determined by consumer valuations but by whatever

1. A linear trend fit over 1970–1973 implies 1974 shortfalls of 8–9% in both EIA petroleum products supplied (U.S. Energy Information Administration 2025) and FHWA net gasoline volume taxed (Federal Highway Administration 2012).

breaks the tie: small differences in transportation costs, regulatory discretion, historical consumption patterns.

The formal consequence is what we call the Chaos Theorem. Since revenue is equalized, any cost advantage, however small, determines the entire allocation. The supplier's problem reduces to cost minimization over the feasible set. Per-unit costs are linear in quantity, the feasible set (bounded by supply and market demands) is a convex polytope, and a linear objective over a convex set achieves its minimum at a vertex. Markets are filled in order of increasing cost until supply is exhausted; served markets receive their full demand at the controlled price while the remaining markets split whatever is left. The "chaos" is in how the vertex is selected. When two markets have nearly equal delivery costs, the economy sits near a knife-edge. A refinery outage, a pipeline repair, a regulatory tweak—any small change in relative costs can flip which market gets filled and which gets starved. The allocation jumps discontinuously even though the parameter change is vanishingly small. Theorem 2 formalizes this: for every neighborhood of a knife-edge parameter, there exist paths along which welfare jumps discretely in both directions.

This corner-solution structure differs from benchmarks in prior work. The standard Harberger triangle is the *minimum* welfare loss compatible with a binding ceiling; it requires the very price variation a ceiling forbids. Corners are worse because they maximize the shadow price gap: starved markets have desperate consumers willing to pay far above the ceiling, while oversupplied markets clear at the controlled price. Each unit sent to the wrong market destroys value equal to this gap. Glaeser and Luttmer (2003) quantify misallocation under rent control assuming random allocation, where every apartment has equal probability of reaching any tenant. Random allocation is an interior outcome with its own smoothness: scarcity is spread widely and small parameter changes produce small welfare changes. We show that equilibrium generically occurs at neither the

Harberger nor Glaeser-Luttmer benchmark. Cost-minimizing suppliers drive allocations to vertices, not interiors (as in Murphy, Shleifer, and Vishny (1992)'s analysis of partial reform). Corners are not an assumption but an *outcome* about what cost-minimizing suppliers choose. The correct benchmark is corners, not random, and corners generate qualitatively different welfare properties: losses far larger than either efficient or random distributions, and discontinuous jumps from small parameter perturbations.

The mechanism is not specific to gasoline or to geography. The term "submarket" covers any dimension along which a market is segmented: space, time, product mix, stages of production. Whenever a ceiling prevents prices from varying across segments, sellers become indifferent, any friction selects a corner, and small shifts in friction flip the system between corners. Set a ceiling at the summer price, and you eliminate the incentive to store oil for winter. Cap refined-product prices and minute differences in refining costs redirect crude toward gasoline and away from diesel. Control output prices while leaving input costs free, and you break the price linkages that coordinate production stages. For example, in the summer of 1973, frozen chicken prices with uncontrolled feed costs led farmers to destroy over a million baby chicks rather than raise them at a loss. "It's cheaper to drown 'em than to put 'em down and raise 'em," one Texas farmer explained (Associated Press 1973). At economy-wide scale, the Chaos Theorem provides micro-foundations for the shortage economy that Kornai (1992) documented in socialist systems: pervasive price suppression generically produces corner allocations that shift unpredictably because they depend on economically irrelevant parameters.

Figure 1 documents this dispersion for the peak of the crisis at the state level, showing the percentage of stations rationing fuel (either completely out of fuel or limiting purchases). Of sampled stations, Connecticut and Massachusetts exceeded 90 percent, while Idaho, Montana, Utah, Hawaii and Wyoming reported no problems. A natural interpretation is that Idaho got lucky and Connecticut got unlucky. But think about what efficient allocation would look like: with a 9 percent national shortfall, every market should experience roughly a 9 percent reduction, or at least *some* reduction. Zero rationing means those states received more than their efficient share. The states with abundant fuel are a sign of the *problem*. Under efficient allocation, a 9 percent national shortfall would translate into a roughly proportional quantity reduction everywhere. The states with zero rationing were, in a sense, overserved, since they had enough to satisfy demand even at the controlled price, not just at the market price. At market prices, vacationers in rural areas would never have outbid commuters in cities for scarce fuel. Under a binding ceiling, both pay the same price, so there is no mechanism to redirect supply. The rationing in Connecticut and the abundance in Idaho are two sides of the same misallocation.

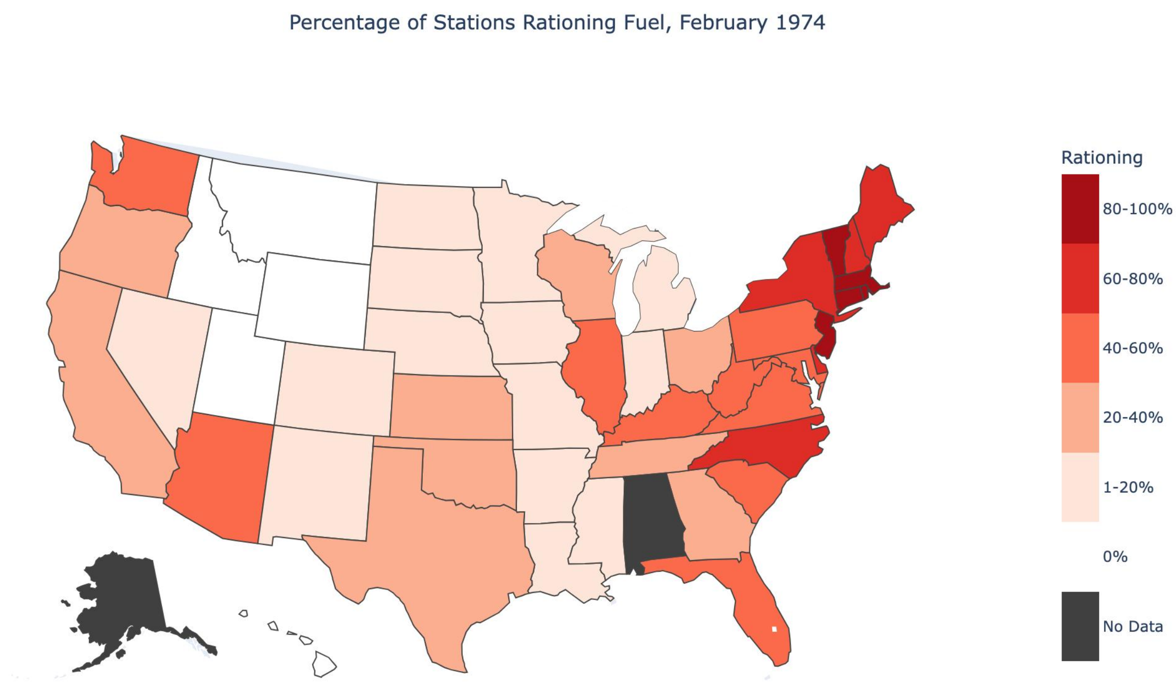

Figure 1: Percentage of gasoline stations rationing fuel by state, February 1974. Rationing includes stations completely out of fuel and those limiting purchases (e.g., maximum gallons per customer). Data are from AAA surveys of sampled stations (AAA 1974).

This reframes the puzzle. The question is not why some states rationed. With a 9 percent national shortfall, rationing was inevitable. The question is why other states exhibited no rationing. Why should some markets clear while others run dry? The answer is that efficient allocation requires price variation, which a binding ceiling explicitly forbids. Once prices cannot adjust, there is no mechanism to direct scarce supplies toward their highest-value uses. The patchwork pattern of Figure 1 is not an accident. It is the generic outcome when prices cannot signal scarcity.[2]

The Chaos Theorem highlights an identification problem. Under a ceiling, equilibrium allocations are generically corners: some markets clear at the controlled price while others are starved. Welfare depends critically on valuing units at these extreme allocations. But demand is identified only locally, near observed price-quantity pairs. Computing welfare at a corner requires integrating inverse demand deep into a region no estimation has reached. Parametric approaches settle that extrapolation by assuming functional form, perhaps bolstered with elasticity sensitivities. We instead ask what welfare statements are possible without committing to a specific extrapolation.

We prove sharp bounds on misallocation losses requiring no parametric demand assumptions. The key structural insight is that efficient allocation equalizes shadow prices across markets, so all markets share a common shadow price p^*. The adding-up constraint $\sum_i q_i(p^*) = \bar{Q}$ converts the infinite-dimensional problem over demand functions into a one-dimensional search over the scalar p^*. Within each market, the extremal demand curves that attain the bounds are piecewise linear, with slopes at the elasticity-implied endpoints. Instead of being a robustness check, robust bounds are the methodologically appropriate response to a challenge intrinsic to the theoretical problem.

How large are misallocation losses? Using station-level survey data from the AAA presented to President Ford during the crisis, we find exactly the corner-solution struc-

2. Pure corner solutions are themselves an extreme case. They assume away contractual rigidities, inventory buffers, and other frictions that slow adjustment. But as a limiting benchmark, corners have empirical bite.

ture the Chaos Theorem predicts: 62.3 percent of stations were operating normally, 27.6 percent were limiting purchases, and 10.1 percent had run out of fuel entirely. This is not a smooth 9 percent reduction spread evenly across stations. Open stations satisfied their customers even at the low, controlled price; closed and limiting stations received the residual. The shadow-price gap is large: in the upper-bound configuration, consumers at rationed stations valued gasoline at several multiples of the controlled price, an arbitrage opportunity that price controls created but prevented from closing. Our robust interval implies a misallocation-to-Harberger ratio of roughly 1 to 9. Even at the lower bound, misallocation roughly equals the quantity-reduction loss; at the upper bound, the Harberger triangle accounts for barely one-tenth of total welfare cost. The quantity reduction is not the main cost of price controls. The misallocation is.

The remainder of the paper proceeds as follows. Section 2 reviews related literature. Section 3 introduces a two-market example. Section 4 presents the general model. Section 5 develops the Chaos Theorem, and Section 5.4 illustrates the chaos mechanism with simulations. Section 7 applies our framework to the 1973–74 gasoline crisis, developing robust bounds on misallocation losses. We defer proofs to Appendix A.

2 Literature Review

Our paper connects three distinct strands of research: the classic welfare-economics treatment of price controls, the modern literature on misallocation, and recent work that uses robust or worst-case methods to bound welfare when key allocation details are unknown.

The canonical approach traces back to Harberger (1954), who measured the efficiency cost of quantity reductions with the familiar ”triangle.” Early analyses of price ceilings took the triangle as their primary yardstick. (e.g. Oi (1976), Gordon (1973), and Rockoff (1984) for U.S. wartime experience).

A subsequent literature pointed out that the triangle misses the welfare costs of non-price rationing. Barzel (1974) showed that when money prices are suppressed, rationing by waiting emerges: consumers queue until the shadow "time price" clears the market. Frech and Lee (1987) and Deacon and Sonstelie (1985, 1989) applied this framework to the 1973–74 gasoline crisis, finding that queuing costs alone could exceed the entire rent transfer from producers to consumers. Frech and Lee (1987) explicitly consider cross-market allocation, analyzing inefficiencies between urban and rural gasoline markets in California. In their framework, optimal rationing by queuing would equalize *time prices* (money price plus waiting cost) across markets. Misallocation occurs when time prices differ: urban consumers faced longer queues than rural consumers, generating welfare losses from the failure to equalize shadow prices. But crucially, their analysis assumes an interior solution: every market receives positive supply, and inefficiency arises from *unequal* shadow prices across locations. Our paper identifies a distinct source of welfare loss that this queuing literature cannot capture. The Chaos Theorem predicts corner solutions where some markets receive literally nothing, an allocative loss that exists even in a hypothetical world with zero queuing costs.

Our paper is most closely related to Davis and Kilian (2011), who study misallocation in the market for natural gas. Brunt (2025) makes a similar pedagogical point, arguing that misallocation losses under price controls are often larger than the quantity-reduction triangle. Murphy, Shleifer, and Vishny (1992) analyze misallocation under partial price liberalization, showing that when some markets are freed while others remain controlled, supply flows to uncontrolled markets until controlled markets are starved. We extend their framework by characterizing the generic structure of equilibrium allocations (corners) and the discontinuous dependence on nuisance parameters.

A foundational question is whether prices or rationing better allocate scarce goods. Weitzman (1977) shows that the answer depends on the joint distribution of needs and income: the price system dominates when wants are dispersed or incomes are equal, while rationing can dominate when needs are uniform but incomes are unequal. The mechanism-design literature formalizes this tradeoff: Condorelli (2013) characterizes when market mechanisms versus non-market mechanisms are optimal, while Akbarpour, Dworczak, and Kominers (2024) show that non-market allocation dominates when willingness to pay is uninformative about the designer's welfare weights. Kominers and Dworczak (2025) define price gouging as occurring when lowering the price from the market-clearing level would raise utilitarian welfare. Our paper abstracts from redistributive motives, focusing instead on the allocative inefficiency that arises when a binding ceiling segments markets and blocks arbitrage.

Our paper builds on the literature studying random allocation under binding price ceilings. Glaeser and Luttmer (2003) show that apartments do not necessarily flow to their highest-value users under rent control, and they quantify the resulting misallocation under random allocation. Bulow and Klemperer (2012) characterize when consumer surplus rises or falls under different demand-curve shapes.

We focus instead on market segmentation: markets are naturally segmented by geography, by transaction and transport costs, by final use, and by position in the structure of production. We show that equilibrium generically occurs at neither the Harberger benchmark (efficient allocation) nor the Glaeser-Luttmer benchmark (random allocation). Cost-minimizing suppliers drive allocations to vertices, not interiors.

Recent empirical work on rent control, such as Diamond, McQuade, and Qian (2019) on San Francisco or the comprehensive survey by Kholodilin (2024), documents substantial misallocation in housing, though identification of the allocative component remains challenging. Buurma-Olsen et al. (2025) make progress on this front, finding that Dutch public-housing tenants consume units differing 7.5 percent from their efficient allocation. We provide a general theoretical framework that nests these settings as special cases.

Our robust bounds draw on the literature measuring welfare without fully specifying demand. Hausman and Newey (1995) estimate consumer surplus using nonparametric methods; Hausman and Newey (2016) show that with unobserved heterogeneity, welfare is only partially identified. Kremer and Snyder (2018) characterize worst-case deadweight loss, showing it can approach the entire surplus under extreme demand shapes. Kang and Vasserman (2025) show that familiar functional forms are extremal within canonical demand families; Kang (2026) extends the robustness approach to the prices-versus-quantities choice. We contribute by applying these methods to misallocation losses across segmented markets, and by showing that the bound-finding problem reduces to a one-dimensional optimization with piecewise-linear extremals.

3 Two Market Example

The simplest setting that exhibits misallocation requires just two submarkets. Consider a national market divided into two submarkets with demand functions $D_1(p)$ and $D_2(p)$. A binding price ceiling $\bar{p}$ creates a total shortage of $D(\bar{p}) - S(\bar{p})$, where $S(p)$ is aggregate supply. We denote total supply under the ceiling as $\bar{Q} \equiv S(\bar{p})$. The quantities allocated must sum to this total: $q_1 + q_2 = \bar{Q}$. No market can receive more than it demands at the ceiling price, otherwise the market price would fall below the ceiling. This gives upper bounds $q_1 \leq D_1(\bar{p}) \equiv \bar{q}_1$ and $q_2 \leq D_2(\bar{p}) \equiv \bar{q}_2$. Together, these constraints define the feasible set.

Figure 2 illustrates the welfare consequences of different allocations. In Panel A, the shortage is allocated efficiently: shadow prices equalize across markets at p^*, yielding deadweight losses a and b that sum to the national triangle c. This is the Harberger case. Panel B shows a corner allocation where Market 1 receives its full demand at the controlled price and Market 2 gets whatever is left. The welfare loss is nearly an order of magnitude larger. The Harberger triangle is not an average outcome but the *smallest* welfare cost compatible with a binding ceiling.

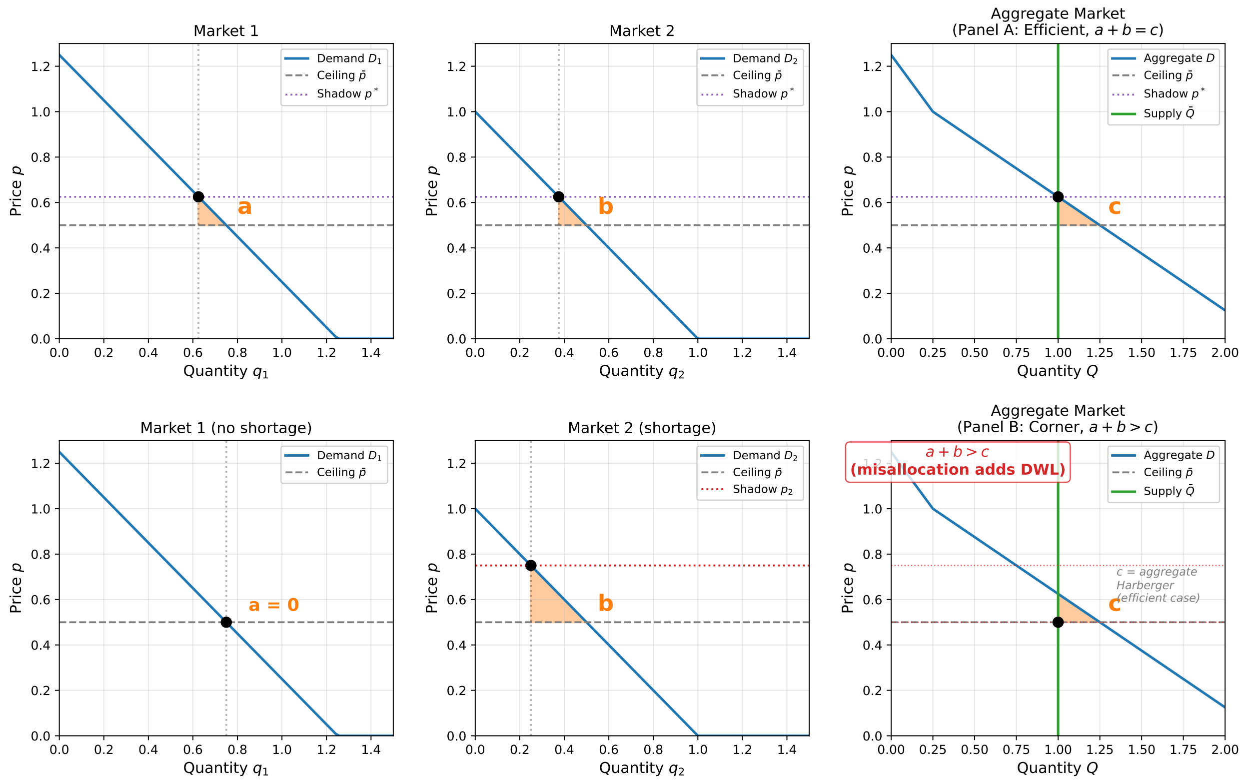

Figure 2: **Two-Market Misallocation under Price Ceiling.** Panel A: efficient allocation with equalized shadow prices; $a + b = c$. Panel B: corner allocation where Market 1 receives full demand; shadow prices diverge and $a + b > c$.

Why should we expect corner allocations rather than the efficient interior? Figure 3 shows the geometry. The constraint $q_1 + q_2 = \bar{Q}$ defines a downward-sloping line; the upper bounds $q_i \leq \bar{q}_i$ truncate it to a line segment. The endpoints E_1 and E_2 are corner solutions.

Between E_1 and E_2 lies a continuum of feasible allocations. Without price controls, arbitrage would push the economy toward the efficient interior. But a binding ceiling eliminates this incentive: with money prices frozen at $\bar{p}$ everywhere, sellers are indifferent across all feasible allocations. As argued in the introduction, this indifference is resolved

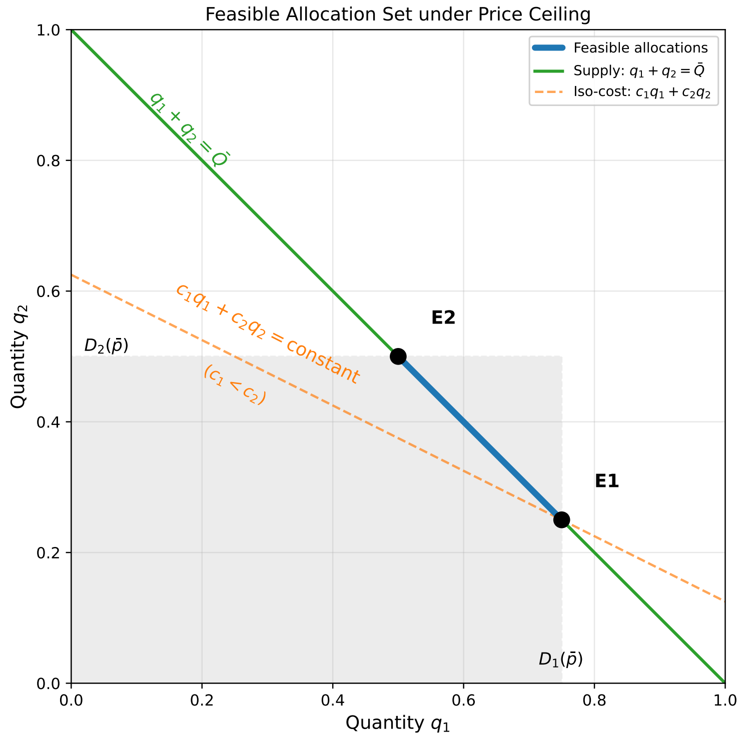

Figure 3: The feasible set is a line segment. Corners E_1 and E_2 are the only generic equilibria under price controls.

by cost minimization. Suppose transporting gasoline to Market i costs c_i per unit. If $c_1 < c_2$, supply flows entirely to Market 1; if $c_2 < c_1$, it flows entirely to Market 2. The continuum of feasible allocations collapses to a single corner.

When c_1 and c_2 are nearly equal, the economy sits near a knife-edge. Any small change in relative costs can flip the entire allocation from E_1 to E_2. Because welfare varies along the feasible set, these discontinuous allocation jumps produce discontinuous welfare jumps. We formalize this in Section 5.

4 General Model

We now generalize the two-market example to an arbitrary number of submarkets with general demand functions. The goal is to characterize, for any market structure and any binding price ceiling, the full range of welfare outcomes, from the efficient allocation that minimizes deadweight loss to the worst-case allocation that maximizes it.

The term ’submarket’ is deliberately flexible. Submarkets may be geographic—states, cities, countries—but can also span temporal segments, end uses, stages of production, and quality tiers. Likewise, ’transport cost’ means any cost of moving product between segments. Set a ceiling at the "summer price" and you eliminate the incentive to ’transport’ oil to winter via storage. Cap refined-product prices and minute differences in refining costs can redirect nearly all crude toward gasoline, starving diesel and jet fuel. Segments are manifold but the mechanism is the same: price controls induce seller indifference across allocations so any friction, however small, selects a corner; and small shifts in friction flip the system between corners.

Formally, a binding ceiling $\bar{p}$ fixes the aggregate quantity at $\bar{Q} = S(\bar{p})$, where S is the aggregate supply curve. Let D denote aggregate demand. We assume that $\bar{p}$ is binding in the sense that $D(\bar{p}) > S(\bar{p})$. There are $n \geq 1$ segmented submarkets. In market i, inverse demand $P_i(\cdot)$ is continuous and nonincreasing on $[0, D_i(\bar{p})]$.

We formalize market segmentation through the concept of a decomposition: a partition of aggregate demand into submarket demands. A decomposition specifies how many consumers in each submarket are willing to buy at any given price. The inverse demand $P_i(q_i)$ in submarket i, the price at which exactly q_i units would be demanded, represents the marginal willingness to pay, or *shadow price*, when quantity q_i is allocated to that market.

Definition 1. A collection of direct demand functions $\{D_i\}_{i=1}^{n}$ with D_i continuous and nonincreasing in price is a *decomposition* if $D(p) = \sum_i D_i(p)$ on $[0, \infty)$.

The price ceiling $\bar{p}$ constrains feasible allocations in two ways. First, total quantity cannot exceed supply: $\sum_i q_i = \bar{Q} \equiv S(\bar{p})$. Second, no submarket can receive more than its demand at the ceiling price, or the local price would fall below the ceiling. Define

$\bar{q}_i \equiv D_i(\bar{p})$ as the maximum quantity submarket i can absorb at price $\bar{p}$. The feasible set is then:

$$\mathcal{F} := \left\{ q \in \mathbb{R}^n_+ : \ 0 \le q_i \le \bar{q}_i \ \forall\, i, \ \sum_{i=1}^n q_i = \bar{Q} \right\}.$$

This is a convex polytope (Grünbaum 2003), meaning the intersection of a hyperplane (the adding-up constraint) with a box (the upper and lower bounds). Its extreme points, or vertices, correspond to allocations where all but one submarket is either completely served ($q_i = \bar{q}_i$) or completely starved ($q_i = 0$).

To measure welfare, we focus on consumer surplus. In submarket i, consumer surplus at allocation q_i, measured relative to the ceiling price $\bar{p}$, is the integral of willingness to pay above the price paid:

$$W_i(q_i) := \int_0^{q_i} [P_i(x) - \bar{p}]\, dx.$$

Aggregating across submarkets gives total consumer surplus: $W(q) := \sum_{i=1}^n W_i(q_i)$.

It will sometimes be convenient to work with *gross* surplus, or the area under the inverse demand curve without subtracting the price paid:

$$\widetilde{W}_i(q_i) := \int_0^{q_i} P_i(x)\, dx, \qquad \widetilde{W}(q) := \sum_{i=1}^n \widetilde{W}_i(q_i).$$

Since $W(q) = \widetilde{W}(q) - \bar{p}\,\bar{Q}$ and total quantity $\bar{Q}$ is fixed on $\mathcal{F}$, maximizing consumer surplus W is equivalent to maximizing gross surplus $\widetilde{W}$. This equivalence simplifies the analysis: we can ignore the constant $\bar{p}\bar{Q}$ term and focus on gross surplus alone.

What allocation maximizes consumer surplus? Under standard conditions, surplus is maximized when the marginal value of the good is equalized across all submarkets that receive positive supply. If one market has a higher shadow price than another, we could increase total surplus by reallocating a unit from the low-value market to the high-value one. This intuition guides the following definition.

Definition 2. An allocation $q \in \mathcal{F}$ *equalizes shadow prices* if there exists $p^* \geq \bar{p}$ such that for every i,

I. $0 < q_i < \bar{q}_i \Longrightarrow P_i(q_i) = p^*$,

II. $q_i = \bar{q}_i \Longrightarrow P_i(q_i) \geq p^*$, and

III. if $q_i = 0$, then $P_i(0) \leq p^*$.

Proposition 1. *q^* maximizes W over $\mathcal{F}$ if and only if q^* equalizes shadow prices.*

The proof is standard: if shadow prices differ, shifting a unit from the low-shadow-price market to the high-shadow-price market increases total surplus by the price gap, contradicting optimality.

Letting $q^* \in \mathcal{F}$ denote an equal-shadow-price allocation from Proposition 1, we now note the following decomposition.

Definition 3. Let $q^* \in \mathcal{F}$ denote an equal-shadow-price allocation from Proposition 1. The *misallocation deadweight loss* of a feasible allocation $q \in \mathcal{F}$ is the welfare shortfall relative to q^*:

$$\mathcal{L}_{Mis}(q) \;:=\; W(q^*) - W(q) \;=\; \sum_{i=1}^{n} \left[W_i(q_i^*) - W_i(q_i) \right].$$

This definition isolates the welfare loss from *how* the fixed total $\bar{Q}$ is distributed across submarkets. It holds constant the aggregate quantity reduction (the source of the Harberger triangle) and asks: given that we must allocate exactly $\bar{Q}$ units, how much additional surplus is destroyed by allocating them inefficiently?

To build intuition, consider the expression in integral form:

$$\mathcal{L}_{Mis}(q) = \sum_{i=1}^{n} \int_{q_i}^{q_i^*} \left[P_i(x) - \bar{p} \right] dx \;=\; \sum_{i=1}^{n} \int_{q_i}^{q_i^*} P_i(x)\, dx,$$

where the second equality uses $\sum_i (q_i^* - q_i) = 0$ to cancel the $\bar{p}$ terms. Each integral measures the value of the units that submarket i receives under the efficient allocation but not under q (if $q_i < q_i^*$), or the value of excess units that i receives under q but not under the efficient allocation (if $q_i > q_i^*$). Markets that are under-served lose high-value units; markets that are over-served gain low-value units. The difference is pure waste.

The misallocation deadweight loss has an important economic interpretation: it represents the value that would be created if arbitrageurs could freely reallocate the fixed supply $\bar{Q}$ across markets. Under normal market conditions, price differences incentivize such reallocation. But when prices are frozen at $\bar{p}$, these incentives vanish. The welfare loss is precisely the forgone gains from trade. They are the dollars left on the table when goods do not flow to their highest-value uses.

4.1 Worst-Case Allocations

We have established that equalizing shadow prices across submarkets minimizes deadweight loss for any fixed aggregate quantity $\bar{Q}$. We now characterize the opposite extreme: which allocation *maximizes* deadweight loss?

Recall that

$$V(q) \equiv \mathcal{L}_{Mis}(q) = \underbrace{\sum_i \int_0^{q_i^*} P_i(x)\,dx}_{\text{constant in } q} - \sum_i \int_0^{q_i} P_i(x)\,dx = W(q^*) - W(q).$$

As we noted earlier, maximizing misallocation $V(q)$ is equivalent to minimizing the gross consumer surplus $\widetilde{W}(q) := \sum_i \int_0^{q_i} P_i(x)\,dx$ over the feasible set $\mathcal{F}$.

Intuitively, the worst case concentrates supply in markets where it generates the least value. Markets with high choke prices receive nothing; markets with low shadow prices are flooded.

Theorem 1. *Assume each P_i is continuous and weakly decreasing on $[0, \bar{q}_i]$ and $\mathcal{F} \neq \emptyset$. Then,*

(i) V is maximized at an extreme point of $\mathcal{F}$, where $q_i \in \{0, \bar{q}_i\}$ for all but at most one i.

(ii) For any worst-case allocation $q^w \in \mathcal{F}$, there exists a cutoff $\lambda_{\text{worst}} \in [0, \infty)$ such that for every market i,

$$\begin{aligned}
&\textit{if} \quad 0 < q_i^w < \bar{q}_i, && \textit{then } P_i(q_i^w) = \lambda_{\text{worst}}, \\
&\textit{if} \quad q_i^w = 0, && \textit{then } P_i(0) \;\geq\; \lambda_{\text{worst}}, \\
&\textit{if} \quad q_i^w = \bar{q}_i, && \textit{then } P_i(\bar{q}_i) \leq \lambda_{\text{worst}}.
\end{aligned}$$

(iii) If $\bar{Q} \leq \bar{q}_j$ for some j, then among single-market vertices $(\bar{Q}, 0, \ldots, 0)$ the worst case is any

$$j \in \arg\min_i \int_0^{\bar{Q}} P_i(x)\, dx\,,$$

that is, the market with the smallest average value up to $\bar{Q}$ (not necessarily the smallest marginal price $P_i(\bar{Q})$).

We defer the formal proof of this theorem to Appendix A.1. Part (ii) provides necessary KKT-type conditions for any worst-case allocation; identifying the global worst among candidate vertices requires comparing average values as in part (iii). Overall, the worst outcome is a form of all-or-nothing. Markets with choke prices above the threshold receive nothing; those below it receive everything. This corner structure maximizes the shadow price gap, hence the welfare loss.

5 The Chaos of Price Controls

The previous section characterized the range of possible welfare outcomes: from the best case (equalizing shadow prices) to the worst case (concentrating supply in low-value markets). But this leaves open a crucial question: *which* allocation should we expect to observe in practice?

Under market clearing, arbitrage answers this question: suppliers ship to wherever prices are highest, and prices rise in scarce markets until shadow prices equalize. Under price controls, revenue per unit is frozen at $\bar{p}$ everywhere, so profit maximization reduces to cost minimization. The market no longer selects allocations; costs do. Whatever breaks the supplier's indifference, no matter how small, determines the entire allocation.

Formally, we show small perturbations in underlying parameters can cause discontinuous jumps in allocations and welfare under price controls. The key insight is that once price variation is removed, the equilibrium allocation becomes determined by arbitrarily small differences in costs or capacities.

5.1 Setup

Let $\theta \in \Theta \subset \mathbb{R}^m$ be a compact parameter space governing submarket inverse demands $P_i(x;\theta)$, unit costs $c_i(\theta)$, and total quantity $\bar{Q}(\theta) > 0$. We assume (i) *Demand regularity*: for each i and every θ, the map $x \mapsto P_i(x;\theta)$ is continuous and strictly decreasing; (ii) *Parameter continuity*: for each x, the map $\theta \mapsto P_i(x;\theta)$ is continuous. We also assume (iii) *Cost continuity*: submarket unit costs $c_i(\theta)$ and the total quantity $\bar{Q}(\theta) > 0$ are continuous in θ.

Without price controls, the unique welfare-maximizing allocation $q^*(\theta)$ varies continuously with θ by Berge's Maximum Theorem. Small changes in parameters produce small changes in allocations and welfare.

5.2 Discontinuity Under Price Controls

Now impose a binding ceiling $\bar{p}$. Define $\bar{q}_i(\theta) := D_i(\bar{p};\theta)$ and $\bar{Q}\,(\theta) := S(\bar{p};\theta)$. Under price controls, profit-maximizing suppliers minimize delivery costs since all units sell at the same price $\bar{p}$. Thus, the allocation is determined by cost minimization:

$$\min_{q \in \mathcal{F}(\theta)} c(\theta) \cdot q,$$

where $\mathcal{F}(\theta) = \{q \in \mathbb{R}^n_+ : \sum_{i=1}^n q_i = \bar{Q}(\theta), 0 \leq q_i \leq \bar{q}_i(\theta)\}$.

Under generic conditions (distinct costs across markets, and total quantity not exactly equal to any subset sum of capacities), the optimizer has a stark structure: markets are filled in order of increasing cost, with at most one market partially filled. Crucially, when two markets have nearly equal costs, an infinitesimal parameter change can flip the entire allocation between them.

5.3 Chaos Result

Let two feasible allocations v and w differ only in which markets receive supply–v fills certain low-cost markets while w fills others. Specifically, suppose allocation w gives market s more units than v does ($w_s > v_s$), while market r receives fewer units ($w_r < v_r$), with the same total quantity. Define the welfare jump–which, here, equals the difference in gross surplus–between them as

$$\Delta W = \int_{v_s}^{w_s} P_s(x)dx - \int_{w_r}^{v_r} P_r(x)dx,$$

where r and s are the markets whose allocations differ.

The theorem formalizes the knife-edge intuition: when costs are nearly equal, infinitesimal parameter changes flip which market is served and which is starved, producing discrete welfare jumps from continuous parameter paths.

Theorem 2 (Chaos). *Suppose costs $c_i(\theta)$ are such that two adjacent vertex allocations (vertices of $\mathcal{F}$ differing in exactly two coordinates) v and w are both optimal at some parameter θ^*. If*

$\Delta W \neq 0$ (a generic condition), then for every neighborhood U of θ^* there exist smooth parameter paths through U along which

(i) The optimal allocation jumps discontinuously between v and w;

(ii) Welfare jumps by $+|\Delta W|$ along one path and $-|\Delta W|$ along another.

In particular, welfare is not locally monotone in parameters: arbitrarily small perturbations can move welfare up or down by discrete amounts.

Corollary 1. *For every neighborhood U of θ^*, there exist parameters $\theta^-, \theta^+ \in U$ with $\|q^*(\theta^+) - q^*(\theta^-)\|_1 = 2\Delta > 0$, where Δ is the quantity shifted between markets.*

The proofs, which formalize the geometry of normal cones and construct explicit crossing paths, appear in Appendix A.2.

The Chaos Theorem goes beyond standard LP sensitivity. Yes, linear programs with non-unique optima exhibit discontinuous selections, but the economic content here is *why* the supplier's problem becomes linear. Under market clearing, welfare maximization has a strictly concave objective (diminishing marginal valuations), guaranteeing continuous allocations via Berge's Maximum Theorem. Price controls eliminate the price variation that creates this curvature: with revenue equalized, profit becomes linear in quantity. The degeneracy is the generic outcome when the arbitrage mechanism is shut down.

This result explains the patchwork pattern of Figure 1. Under price controls, which markets receive supply depends on cost parameters that may shift unpredictably with weather, refinery maintenance, or regulatory discretion. A small change–say, a pipeline repair that marginally reduces delivery costs to New Jersey–can discontinuously redirect supply away from New York, even though both states' underlying demands are nearly identical.[3]

3. The Chaos Theorem assumes agents are rational cost-minimizers which pushes allocations to corner solutions. In practice, long-term contracts, sticky relationships, or regulatory inertia may prevent imme-

The theorem also explains why which markets are starved is often volatile over time. Week-to-week fluctuations in transportation costs, inventory positions, or bureaucratic priorities could flip allocations between equilibria, generating the erratic patterns documented in contemporary reports.

More broadly, the result shows that price controls introduce a fundamental unpredictability into markets. When prices cannot signal value, allocations become hypersensitive to ”nuisance parameters” that would be irrelevant under market clearing. This hypersensitivity, chaos in the mathematical sense, is not a market failure but a direct consequence of suppressing the price mechanism. Ironically, the chaos of price controls often fuels demands for even tighter controls and increasingly politicized and bureaucratized allocations.

How much smoothing would be required to bring misallocation toward Harberger magnitudes? If half of supply were locked into historical patterns via long-term contracts or regulatory inertia, the remaining half would still face corner incentives, roughly halving misallocation losses but still leaving them several times the triangle. Only near-complete insulation from cost-based reallocation would eliminate the mechanism entirely. The Chaos Theorem is a limiting case, but the qualitative prediction, that misallocation losses dominate quantity-reduction losses, is robust to substantial smoothing.

5.4 Simulation

We illustrate the Chaos Theorem with simulations of 100 gasoline markets arranged on a 10×10 grid. Each city has a per-unit delivery cost drawn uniformly from $[0, 0.1]$. We use a demand specification with a finite choke price, so welfare is well-defined even when

diate jumps to vertices. However, over time, competitive pressures tend to push allocations toward these extreme points, making the chaos result increasingly relevant as the control regime persists. In this sense, the Chaos Theorem is a worst-case outcome.

markets receive zero.[4] Aggregate supply is fixed at $Q = 150$ units, and we impose a price control at 80% of the market-clearing price.

Under price controls, the allocation follows a greedy rule: fill the lowest-cost city to capacity, then the next-lowest, until supply is exhausted. This produces a vertex allocation where served cities receive their full demand while roughly 30 cities receive nothing.

Figure 4 shows allocations under three random cost draws. The top row shows free-market allocations; the bottom row shows price-control allocations for the same costs. Three patterns emerge. First, free-market allocations are nearly identical across scenarios—small cost differences produce small allocation differences. Second, under price controls, many markets receive zero (yellow). Third, under price controls, different cost draws produce radically different allocations. This is the Chaos Theorem: small changes in nuisance parameters dramatically change the final allocation. Misallocation in these scenarios generates a welfare loss of roughly 13%.

Cost structures may also have systematic components: locations close to refineries may have lower costs. These produce different spatial patterns but the same corner-solution logic.

4. We use a Hill (sigmoidal) demand function.

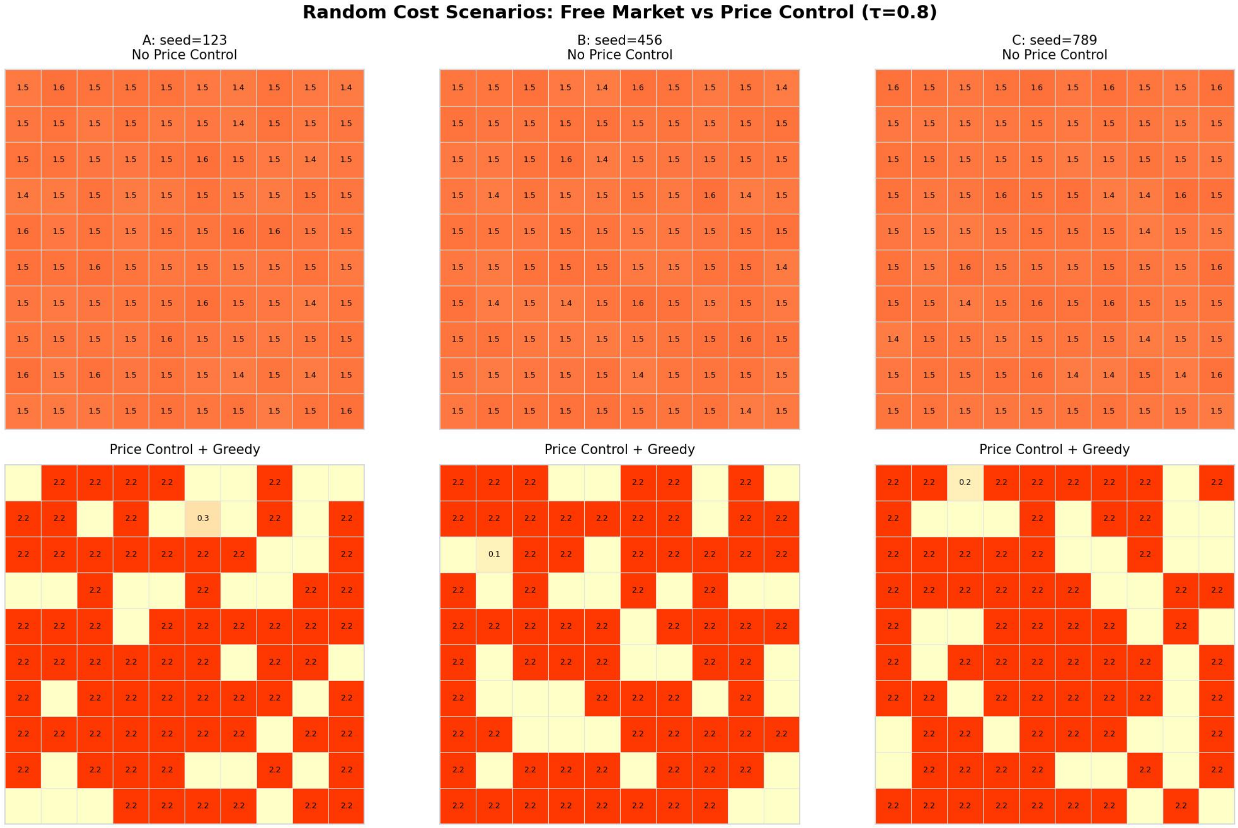

Figure 4: Allocation of a fixed supply ($\bar{Q} = 150$) across 100 cities under free markets (top) versus price controls (bottom). Each column is a different random draw of delivery costs. Under free markets, quantities vary smoothly. Under price controls, low-cost cities fill to capacity while approximately 30 cities go unserved.

6 Robust Bounds on Misallocation

The Chaos Theorem tells us that under price controls we should expect corner allocations, far from market equilibria. We know, rationed stations received only 68% of baseline quantity and often no quantity at all. Elasticity estimates identify demand locally near equilibrium, not at these depressed allocations. Rather than extrapolating via functional form assumption, we develop a partial identification approach. Given only the observed allocation, the ceiling price, bounds on local demand elasticities, and an anchored point on each demand curve (or a bounded interval for that anchor), we derive sharp bounds on misallocation. This is not sensitivity analysis within a parametric family. We do not commit to any functional form. Instead, we ask: across all inverse demand functions consistent with slope bounds, what are the largest and smallest possible welfare losses? The output is an identified interval, not a point estimate, and the interval covers welfare under all robustly possible allocations.

In market i, the unknown inverse demand $P_i\colon \left[0, q_i^{\max}\right] \to \mathbb{R}_+$ is continuous and nonincreasing in quantity q, where each bound $q_i^{\max} \in (0, \infty)$ is an exogenous upper bound on feasible quantities (representing physical capacity, infrastructure limits, or maximum historical consumption), distinct from the demand-derived $\bar{q}_i = D_i(\bar{p})$ used elsewhere. A binding ceiling $\bar{p} > 0$ is imposed. We observe the delivered allocation $q^{\text{obs}} = (q_1^{\text{obs}}, \ldots, q_n^{\text{obs}})$, where

$$\bar{Q} := \sum_{i=1}^{n} q_i^{\text{obs}} \in \left[0, \sum_i q_i^{\max}\right]$$

.

Assumption 1. For all i, P_i is nonincreasing and satisfies

I. $P_i(q_i^{\text{obs}}) = p_{0,i} = \bar{p} - b_i \quad \text{with} \quad b_i \in [\underline{b}_i, \bar{b}_i]$.

II. There exist $g_{i,L} < g_{i,U} < 0$ with $g_{i,L} \leq P_i'(q) \leq g_{i,U}$ for a.e. $q \in \left[0, q_i^{\max}\right]$.

III. $P_i(0) \leq M_i < \infty$

IV. $P_i \in W^{1,\infty}\left(\left[0, q_i^{\max}\right]\right)$, the space of bounded Lipschitz-continuous functions on $\left[0, q_i^{\max}\right]$.

The first three assumptions are the economically meaningful ones, whereas the last one is technical, guaranteeing that the various objects we construct are well-defined. Let $\mathcal{P}_i$ be the set of P_i satisfying Assumption 2 and $\mathcal{P} := \times_{i=1}^{n} \mathcal{P}_i$. We assume $\mathcal{P} \neq \varnothing$.

Let $q_i(\cdot)$ denote the left-continuous generalized inverse of inverse demand $P_i(\cdot)$:

$$q_i(p) := \begin{cases} \inf\left\{x \in \left[0, q_i^{\max}\right] : P_i(x) \le p\right\}, & \text{if the set is nonempty,} \\ q_i^{\max}, & \text{otherwise.} \end{cases}$$

Let $\Phi(P)$ denote the misallocation loss:

$$\Phi(P) := \sum_{i=1}^{n} \int_{q_i^{\text{obs}}}^{q_i^*(P)} \left(P_i(x) - p^*(P)\right) dx,$$

where $q^*(P)$ is the equal-shadow-price split of total quantity $\bar{Q}$ and $p^*(P)$ is its Lagrange multiplier. Define

$$\overline{\Phi} := \sup_{P \in \mathcal{P}} \Phi(P) \qquad \text{and} \qquad \underline{\Phi} := \inf_{P \in \mathcal{P}} \Phi(P).$$

We sketch the result here, as the formal result is quite notation heavy. The fully detailed statement and its proof are given in the appendix.

Within each market, the logic follows Kang and Vasserman (2025): extremal welfare is achieved by demand curves at the slope bounds. Our contribution is recognizing that efficient allocation equalizes shadow prices across markets, so all markets share a common p^*. This adding-up constraint converts the infinite-dimensional problem over demand functions into a one-dimensional search over the scalar p^*.

Proposition 2 (Informal)**.** *Posit an interiority assumption. Then there exist extremizers $P^* \in \mathcal{P}$ attaining $\underline{\Phi}$ and $\overline{\Phi}$ such that each P_i^* is continuous and piecewise affine, with $P_i^{*\prime}(q) \in \{g_{i,L}, g_{i,U}\}$ for a.e. q. Moreover, on the quantity range that matters for the bound (between q_i^{obs} and the equal-shadow allocation), each market can be taken to have at most one kink.*

Computing $\underline{\Phi}$ and $\overline{\Phi}$ reduces to a one-dimensional search over a candidate common shadow price p. For each fixed p, the problem decomposes across markets except for the single adding-up constraint $\sum_{i=1}^{n} q_i(p) = \bar{Q}$.

Proposition 3 is the heart of this robust bounds approach. It turns the bound-finding problem over the huge class $\mathcal{P}$ into a one-dimensional program over $p \in \mathcal{I}$, and shows that the optimizer hits the slope endpoints everywhere except at kinks. The key simplification is that, conditional on p, the objective is *linear* in the inverse-quantity functions $q_i(\cdot)$, so the optimum must lie at the pointwise boundary of the admissible set.

Theorem 3 (Informal)**.** *Posit an interiority assumption. Then, both $\overline{\Phi}$ and $\underline{\Phi}$ are attained by some $P^* \in \mathcal{P}$. Moreover, each P_i^* may be chosen continuous and piecewise affine on $[0, q_i^{\max}]$, with a.e. derivative $P_i^{*\prime}(q) \in \{g_{i,L}, g_{i,U}\}$, and satisfying $P_i^*(0) \leq M_i$.*

The proof appears in Appendix B.

Corollary 2. *Suppose we also optimize over $p_{0,i} \in [\bar{p} - \bar{b}_i, \bar{p} - \underline{b}_i]$. If the candidate shadow price is at least as large as the price implied by the observed quantity in market i; viz., $p \geq P_i(q_i^{\text{obs}})$ and none of the upper or lower bounds on quantities bind over the relevant price range, the upper bound is produced by the largest bias and the lower bound by the smallest bias.*

In all empirical robust-bound results below, we implement the interval-anchor extension of Corollary 2: $p_{0,i}$ is optimized over its admissible interval jointly with the shadow-price search.

Putting these results together, the computation proceeds as follows:

I. For each market i, the slope bounds and anchor $(q_i^{\text{obs}}, p_{0,i})$ pin inverse demand within a cone: any admissible P_i must lie between the steepest and flattest lines through the anchor. Inverting this cone yields the quantity envelopes $\ell_i(p)$ and $u_i(p)$: the smallest and largest quantities market i could absorb at shadow price p under any admissible demand. When a choke cap $P_i(0) \leq M_i$ binds, the upper envelope is trimmed further (Theorem 7).

II. The set of candidate common shadow prices is $\mathcal{I} = \{p : \sum_i \ell_i(p) \le \bar{Q} \le \sum_i u_i(p)\}$, i.e. the prices at which per-market quantity intervals are jointly compatible with the aggregate constraint.

III. For each candidate $p \in \mathcal{I}$, the conditional problem pushes each market to a boundary envelope, but the relevant endpoint depends on both the bound direction (upper versus lower) and whether p lies above or below $p_{0,i}$; equivalently, each market has a boundary target $c_i(p) \in \{\ell_i(p), u_i(p)\}$.

IV. These boundary targets need not satisfy adding-up exactly, i.e. $\sum_i c_i(p)$ need not equal $\bar{Q}$. Feasibility is restored by solving a convex quadratic projection onto $\{x : \sum_i x_i = \bar{Q},\ \ell_i(p) \le x_i \le u_i(p)\}$ that penalizes deviations from $c_i(p)$, with weights $\kappa_i = 1/g_{i,L} - 1/g_{i,U}$. Markets with wider slope uncertainty (larger κ_i) absorb more of this adjustment.

V. The resulting value $\Phi(p)$ is optimized over $p \in \mathcal{I}$ to obtain $\overline{\Phi}$ (or $\underline{\Phi}$). When anchor prices are interval-valued, this becomes a joint optimization over p and the anchor vector p_0 on the admissible hyper-rectangle $\prod_i [p^L_{0,i}, p^U_{0,i}]$ (Corollary 2).

Theorem 7 establishes that the bounds $[\underline{\Phi}, \overline{\Phi}]$ are *sharp*: they are attained by demand functions in the permitted class. This sharpness distinguishes partial identification from sensitivity analysis. Sensitivity analysis with linear demand would vary the slope b and report how misallocation changes, yielding the same numerical interval when the choke-price constraint is non-binding. With a binding $P(0) \le M$, the robust bounds can differ because the extremal piecewise-affine demands can kink to satisfy the choke while steeper linear extrapolations cannot. But sensitivity analysis cannot rule out the possibility that some nonlinear demand (convex, concave, or oscillating within the slope bounds) produces misallocation outside that interval.

The theorem closes this gap. By the fundamental theorem of calculus, any demand function with $P'(q) \in [g_L, g_U]$ must lie within a cone emanating from the observed point. Since misallocation is an integral of the demand curve, it inherits these bounds. The theorem confirms that the extreme misallocation is achieved at the boundary of this cone, by piecewise affine demand with slopes in $\{g_L, g_U\}$, not by some exotic function in the interior. The bounds from linear demand are therefore not conservative approximations; they are exact over the infinite-dimensional class of all Lipschitz functions satisfying the slope constraints.

7 Illustrating Chaos and Misallocation: The 1973–74 Gasoline Crisis

In August of 1971, President Nixon froze all wages and prices in the United States. Many prices were unfrozen after the 90-day scheduled period, but petroleum remained subject to controls with various modifications until August of 1973 when Special Rule No. 1 established mandatory controls for petroleum specifically [5].

Price controls on oil and the sporadic shortages were thus already in place well before the Oil Embargo began in October of 1973. The embargo, however, dramatically increased the market price of oil above the controlled price, transforming what had been moderate distortions into severe shortages. From October of 1973 to January of 1974 the price of oil tripled. Over the first shock, quantity consumed fell by about 9% relative to trend (Yergin 1991).

Quickly following the embargo, the Emergency Petroleum Allocation Act of 1973 (EPAA) ratified President Nixon's earlier executive orders with legislated price controls

5. See Bradley (1997) for a definitive account of the crisis.

and an allocation system based pro-rata on 1972 levels. If total supply fell to 90% of 1972's volume, for example, each buyer would receive 90% of their 1972 allocation. This basic system was then adjusted with numerous exceptions, prioritizing industries such as national defense, essential services, agriculture, and independent refiners, all under an overarching "fair and equitable" guideline (Bradley 1997).

We begin with a formal analysis of the most visible portion of the petroleum allocation under price controls. Namely, the allocation of gasoline across geographies. We then turn to an informal discussion of misallocation across time, sector and other segmentations.

7.1 AAA Data

We digitize station-level survey data collected by the AAA during February 1974. Figure 1 displays the cross-sectional variation in rationing (out-of-fuel closures + purchase limits). The mean out-of-fuel rate by state was 8.2%, but this average masks striking heterogeneity. The survey reports that 10.1% of stations were closed (out of fuel), 27.6% were limiting purchases, and 62.3% were operating normally. We combine this station-level classification with 1972 gasoline station counts and sales data from the U.S. Census of Retail Trade. Our sample includes 48 states: we exclude Alabama, Alaska, and DC.[6]

For the state-level shadow price map (Figure 7), we weight states by 1972 baseline gasoline quantities. The Census reports both dollar sales and gallon sales by state, but gallon data are missing for 12 states. Where gallon sales are observed, we use them directly. Where only dollar sales are observed, we impute gallons as sales divided by an implied price. The implied price is the midpoint of prices from states with observed gallons, excluding Indiana whose implied price is an outlier. Figure 8 in Appendix C shows the same map using only states with observed gallon data (no imputation); the geographic

6. Alabama and Alaska lacked AAA survey coverage.

pattern and welfare bounds are similar. Dropping Indiana from the sample entirely does not change the results.[7]

The AAA data are corroborated by extensive historical accounts of the misallocation chaos. In Maine, for example, heating oil was in such short supply that the state Office of Civil Defense directed each community to prepare emergency shelters (Turkel 2023). In Maryland and Connecticut, lines for gasoline stretched up to five miles long. Yet, *at the same time*, Texas, the Deep South, and the Great Plains were "virtually awash with gasoline" (TIME 1974). Divisions this stark do not occur in an integrated market.[8]

7.2 Welfare Measures

The Harberger deadweight loss measures the welfare cost of the aggregate quantity reduction, assuming the reduced supply is allocated efficiently:

$$\mathcal{L}_{Harb} = \widetilde{W}(q^{base}) - \widetilde{W}(q^{*}) - p_{base} \cdot (Q^{base} - \bar{Q}) \tag{1}$$

Here $\widetilde{W}(q) := \sum_i \int_0^{q_i} P_i(x)\,dx$ denotes gross surplus (the area under demand), in contrast to the net surplus W defined earlier. The final term subtracts baseline expenditure on the missing quantity at p^{base}, yielding the standard Harberger triangle benchmark. We express these welfare losses as percentages of *baseline expenditure*: total consumer spending before the crisis, $p^{base} \times Q^{base}$. With our normalization ($p^{base} = 1$, $Q^{base} = 1$), baseline expenditure equals unity. This metric provides economic interpretation: a misallocation loss of 3.6% means consumers lost welfare equivalent to 3.6% of their pre-crisis gasoline spending. For our calibration, with a 9% aggregate shortage ($\bar{Q} = 0.91$) and $\bar{p} = 0.8$, the

7. Indiana is included in the main specification with its observed gallon data; it is excluded only from the price calibration used to impute missing data for other states.

8. We focus on the 1972-1974 crisis but the same patterns were in evidence in 1979 when an even smaller national quantity reduction $\sim$3.5% led at its height in June of 1979 to the closing of almost every gasoline station in New Jersey, Connecticut and New York City (Gupte 1979).

Harberger triangle is approximately $\mathcal{L}_{Harb} \approx 2.0\%$ of baseline spending.

The misallocation loss $\mathcal{L}_{Mis} = W(q^*) - W(q^{obs})$ captures the additional welfare destruction from inefficient distribution of the reduced supply, as defined in Section 3. Our key statistic is the *Misallocation Ratio*:

$$\mathcal{R} = \frac{\mathcal{L}_{Mis}}{\mathcal{L}_{Harb}} \tag{2}$$

A ratio $\mathcal{R} = 1$ means misallocation doubles the welfare cost relative to a standard Harberger analysis; $\mathcal{R} > 1$ means misallocation is the dominant source of loss. With efficient rationing, $\mathcal{R} = 0$.

7.3 Allocation

We aggregate stations into two markets: open (62.3%) and closed/limiting (37.7%). "Out of fuel" is a point-in-time status from the AAA survey, not a period-average of $q = 0$. Our two-market calibration aggregates closed/limiting stations into a single partially-served submarket. The key economic mechanism is that the controlled price $\bar{p}$ lies below the market-clearing price. At $\bar{p}$, quantity demanded exceeds quantity supplied. Open stations satisfy their customers; closed and limiting stations cannot.

Under a vertex allocation, open stations consume their full demand at the controlled price while closed stations receive the residual. We assume $\bar{p} = 0.8$, consistent with (Frech and Lee 1987). With elasticity bounds $\varepsilon \in [0.2, 0.4]$ (Hughes, Knittel, and Sperling 2008), baseline-anchored linear demand implies

$$q_O \in [1.04,\ 1.08].$$

Closed/limiting stations receive the residual

$$q_C = \frac{\bar{Q} - (1-s) \cdot q_O}{s} \quad \Rightarrow \quad q_C \in [0.63,\, 0.70].$$

For the baseline calibration, we use $\varepsilon = 0.3$, which gives $q_O = 1.06$ and $q_C \approx 0.66$. This allocation satisfies the aggregate constraint and exhibits exactly the corner-solution structure predicted by the Chaos Theorem.

7.4 Station-Level Robust Bounds

Two features of our setting make partial identification essential. First, extreme allocations are where functional form does most of the work. Our welfare calculations integrate inverse demand over quantity ranges extending to $q \approx 0.66$ for closed/limiting stations. Elasticity estimates come from observed price-quantity variation around market equilibrium, identifying the slope of demand locally near $q = 1$, not at $q = 0.66$. To compute welfare at these allocations, we must extrapolate beyond where demand is identified. Linear extrapolation from an elasticity of 0.2 implies $P(0) = 6$; from 0.4, it implies $P(0) = 3.5$. CES demand diverges entirely as $q \to 0$, requiring an arbitrary choke-price normalization. In our calibration, this extrapolation uncertainty moves misallocation estimates by many percentage points of baseline spending (from roughly 2.3% to 18.6% in the pooled no-choke benchmark).

Second, unequal allocations are the generic prediction. The Chaos Theorem tells us to expect extreme dispersion, and the February 1974 data deliver exactly that: 10.1% of stations closed and 27.6% limiting purchases while 62.3% operated normally. This is not a smooth 9% quantity reduction spread uniformly across stations. At the controlled price, open stations satisfy their customers at $q = 1.06$ while closed/limiting stations sell out or ration at $q = 0.66$. This gap is exactly the configuration where parametric welfare calculations are most fragile.

Taken together, the theory predicts allocation dispersion, the data show dispersion, and dispersion is where functional form assumptions bite hardest. We do not want our welfare estimates hostage to how we extrapolate demand into regions no estimation has ever reached. The robust bounds approach sidesteps this problem. We impose only local, transparent restrictions (slope bounds from elasticity estimates and a finite choke price) and optimize over all inverse demands consistent with these assumptions.

The elasticity bounds $\varepsilon \in [0.2, 0.4]$ imply slope bounds $g \in [-5, -2.5]$.[9] To visualize how these bounds translate into welfare uncertainty, we aggregate stations into a pooled two-market benchmark (open versus closed/limiting) where the extremal demand curves can be drawn directly. Figure 5 shows the demand curves consistent with the slope bounds and observed allocation when no choke constraint is imposed. Solid segments are pinned down by the data and slope restrictions; dashed segments are admissible extensions.

Figure 6 adds a choke constraint $P(0) \leq M = 4$. The key intuition is geometric. In the lower-loss panel, the steep branch would extrapolate to $P(0) = 6 > 4$, so the admissible curve must kink and bend to hit the choke cap. In the higher-loss panel, the flat branch extrapolates to $P(0) = 3.5 < 4$, so no kink is required. This mainly trims the closed/limiting upper shadow-price bound at observed quantity (about 2.67 $\rightarrow$ 2.34), lowering the upper welfare endpoint. We therefore present no-choke results as the headline interval and use the choke case as a disciplined robustness check.

The pooled two-market benchmark visualized above gives no-choke bounds of $\mathcal{R} \in$ [1.15, 9.18]. This is a useful diagnostic object because the extremal demand curves can be

9. Throughout, "steep" and "flat" refer to the inverse demand curve $P(Q)$. A steep inverse demand (large $|g|$, low elasticity) means price falls rapidly as quantity rises. This corresponds to a *flat* demand curve $Q(P)$, where quantity responds little to price changes. The figures plot P against Q, so visual steepness matches $|g|$.

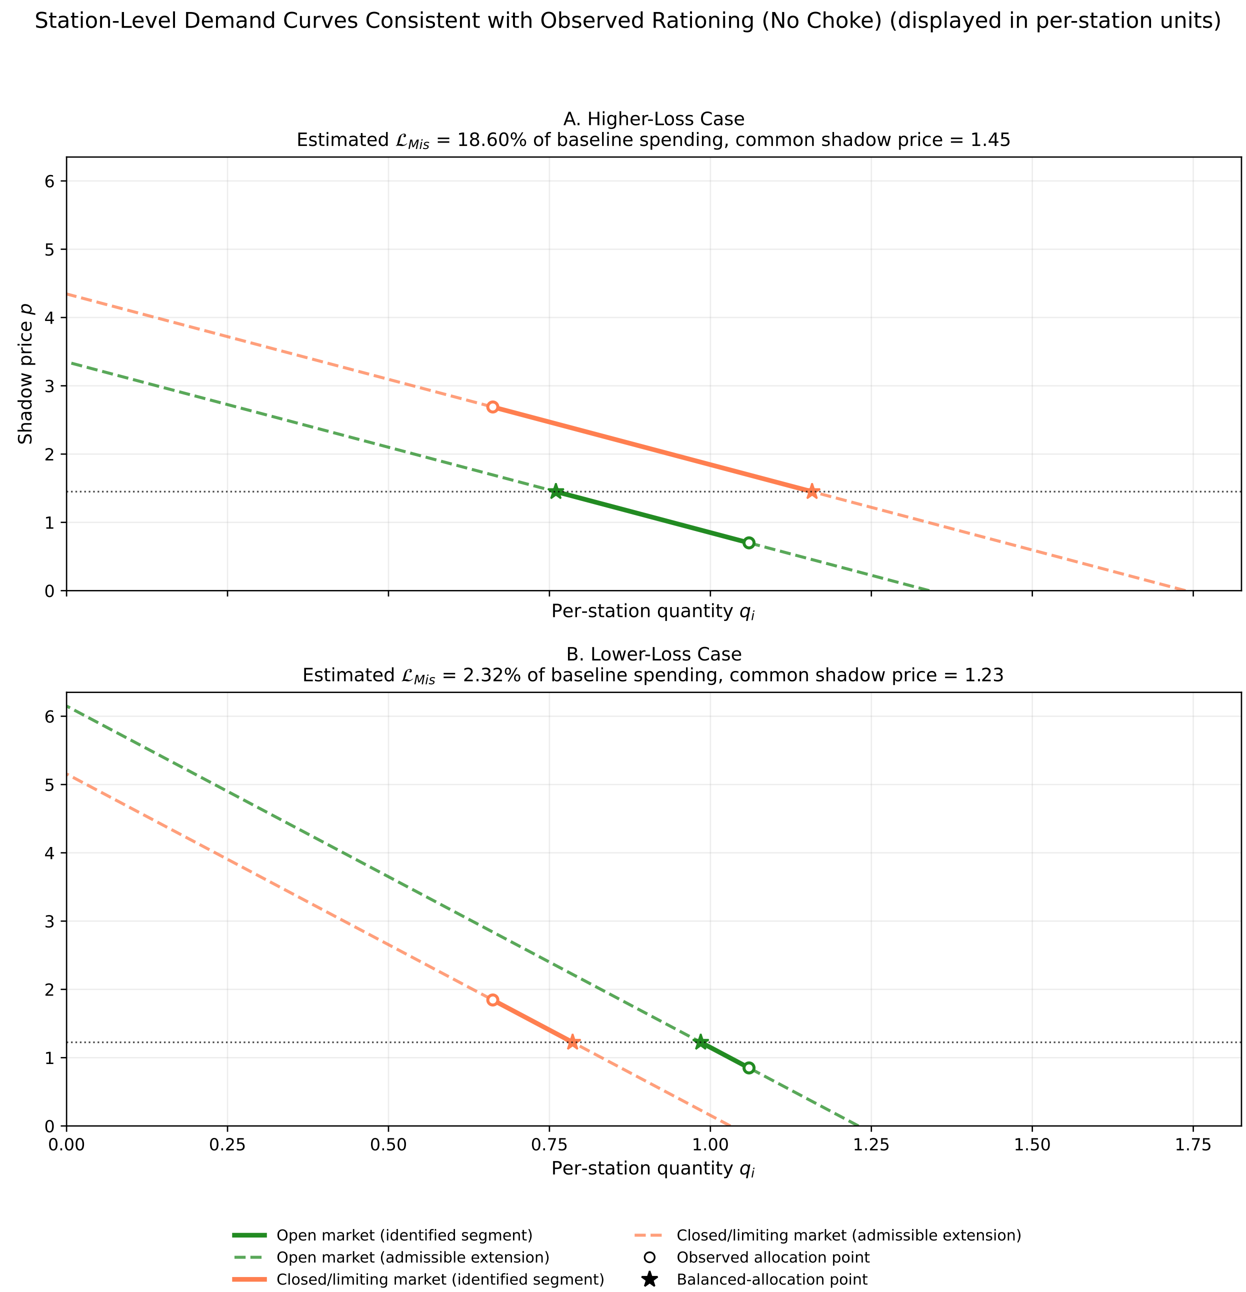

Figure 5: **Station-Level Demand Curves Without Choke.** The two panels show high-loss and low-loss demand configurations consistent with observed open and closed/limiting station quantities when no choke cap is imposed, using the baseline calibration $q_O = 1.06$ and $q_C = 0.66$. Solid segments are identified by the observed allocations and slope bounds; dashed segments are admissible extensions.

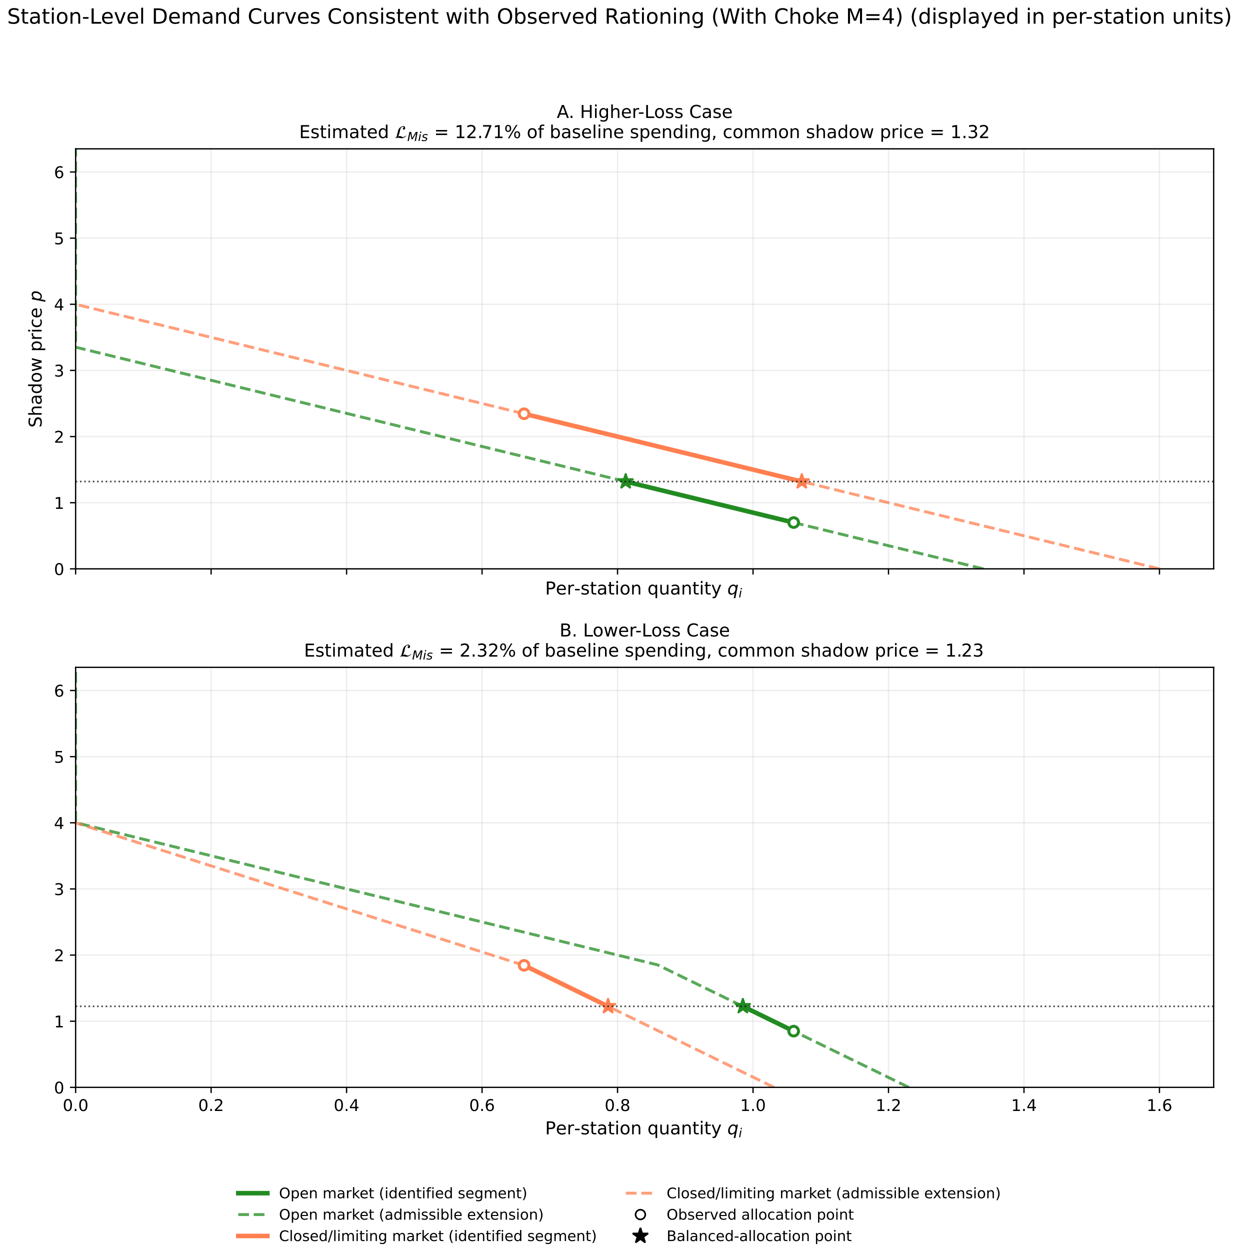

Figure 6: **Station-Level Demand Curves With Choke.** Same construction as Figure 5, using the same baseline calibration and now imposing $P(0) \leq M = 4$. The choke cap narrows admissible shadow-price ranges, which tightens the upper bound on misallocation losses.

visualized directly. Our main empirical specification is different: it disaggregates to the state-by-status level, treating open and non-open stations within each state as separate markets and imposing adding-up jointly across those cells. Each state's weight is proportional to its gasoline consumption (gallons), not its station count, so high-consumption states carry more weight than their station share alone would imply.

7.5 State-Level Shadow Prices

Figure 1 showed substantial geographic variation in rationing severity. We can translate this variation into state-level shadow prices using the state-by-status robust model. For each state i with rationing share r_i (limiting plus out-of-fuel), the state-average shadow price is

$$\bar{p}_{0,i} = (1 - r_i)\, p_{0,i}^{open} + r_i\, p_{0,i}^{nonopen}.$$

Figure 7 maps these state averages for the upper-bound (top) and lower-bound (bottom) configurations. States with high rationing shares (Connecticut, Massachusetts) load heavily on the non-open shadow price and appear with higher state-average shadow prices; low-rationing states (Idaho, Montana, Wyoming) load mostly on the open shadow price and appear lower.

The geographic pattern mirrors Figure 1: where rationing was severe, shadow prices are higher; where rationing was limited, shadow prices are lower. This is the open/non-open mechanism visualized in geographic form.

The economic meaning of these shadow prices is direct. In the upper-bound configuration, the state-average shadow price in Connecticut is 2.5 times the baseline price versus 0.7 times baseline in Montana, where no stations reported shortages. Connecticut consumers valued gasoline at 3.5 times what Montana consumers were paying. Shipping a barrel from Montana to Connecticut would have more than tripled your money: the

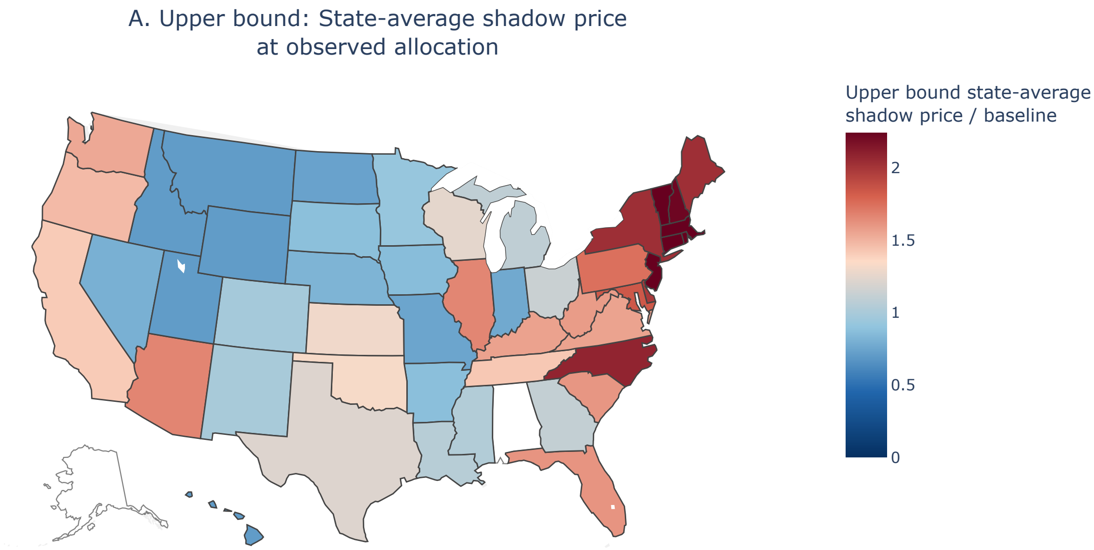

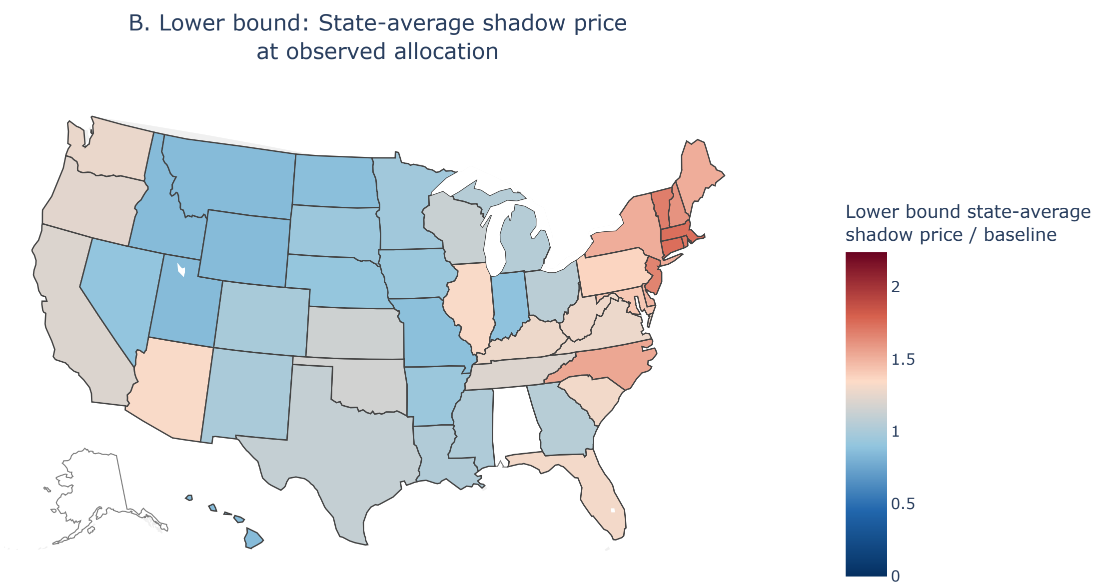

Figure 7: **State-Average Shadow Prices.** Top panel: upper bound. Bottom panel: lower bound. Colors show each state's average shadow price, computed as the rationing-share-weighted average of open and non-open shadow prices from the state-by-status robust bounds.

arbitrage opportunity that price controls created but prevented.

This state-by-status specification uses 91 active state-status markets rather than the pooled two-market benchmark. For each active cell, the anchor price $p_{0,i}$ is allowed to vary over its admissible interval (implied by baseline anchors and slope bounds), and we optimize jointly over these intervals and the common shadow-price condition. Relative to pooled two-market bounds, this additional structure slightly tightens both endpoints:

$$\mathcal{L}_{Mis}/\mathcal{L}_{Harb} \in [1.09, 8.75].$$

Since $\mathcal{L}_{Harb} \approx 2.0\%$ of baseline spending, this ratio implies $\mathcal{L}_{Mis} \in [2.2\%, 17.5\%]$. Even at the lower endpoint, misallocation losses exceed the Harberger triangle; at the upper endpoint, misallocation remains close to an order of magnitude larger.

This slight tightening is reasonable but not mechanical. The state-by-status model adds many cell-level feasibility restrictions that must all be satisfied jointly with a common shadow price and adding-up, which can reduce the set of feasible extremal reallocations relative to a pooled two-market benchmark. At the same time, interval anchor uncertainty in $p_{0,i}$ adds flexibility, so monotonic tightening is not guaranteed ex ante. In our calibration, the state-status primitives are fairly similar across states, so the difference from pooled bounds is modest rather than dramatic.

Table 1 decomposes the bound interval by assumption. Quantities are calibrated once: $q_{open} = 1.06$ and $q_{non\text{-}open} = 0.67$ follow from $\varepsilon_{open} = 0.3$ and the 9% shortage, and are held fixed throughout. What varies is the *welfare elasticity*, the demand slope used to compute the surplus integral.

Row 1 assigns every market the same welfare elasticity. Sweeping it over $[0.2, 0.4]$ changes Φ by a factor of two but leaves $\mathcal{R}$ invariant at 4.4: the slope scales both misallocation loss and the Harberger triangle identically. Row 2 allows each market its own slope within the same range. This heterogeneity is what breaks the invariance, widening $\mathcal{R}$ to [2.5, 4.9]. Allowing interval-anchor uncertainty (row 3) is the main widening step. Adding a choke cap at $M = 4$ (row 4) leaves the lower endpoint unchanged but trims the upper endpoint, because the lower-loss extremal already satisfies the cap while the high-loss extremal is where the choke binds.

Assumptions	$\underline{\Phi}$ (%)	$\overline{\Phi}$ (%)	$\mathcal{R}$
Common $\varepsilon \in [0.2, 0.4]$ (fixed anchors)	4.43	8.86	4.38
Heterogeneous $\varepsilon_i \in [0.2, 0.4]$ (fixed anchors)	4.99	9.96	[2.46, 4.92]
+ Anchor uncertainty ($p_{0,i}$ intervals)	2.21	17.72	[1.09, 8.75]
+ Choke constraint ($M = 4$)	2.21	11.62	[1.09, 5.74]

Table 1: **Assumption-to-Interval Decomposition.** Quantities calibrated at $\varepsilon_{open} = 0.3$ ($q_{open} = 1.06$) and held fixed; welfare elasticity (demand slope for surplus computation) varies as indicated. Φ: misallocation loss as percent of baseline spending. $\mathcal{R} \equiv \mathcal{L}_{Mis}/\mathcal{L}_{Harb}$, where Ψ uses the same ε as Φ in row 1 (matched) and the conservative $\varepsilon = 0.2$ in rows 2–4.

These magnitudes are qualitatively consistent with the queuing equilibria documented by Deacon and Sonstelie (1985). When money prices cannot ration demand, time prices must rise to clear the market. At the height of the crisis, hour-long waits were routine at hard-hit stations, with multi-hour waits common at peak times. Deacon and Sonstelie (1985) estimate the value of waiting time at roughly the after-tax wage rate. For a typical motorist, a two-hour wait represents a time cost comparable to the fill-up itself, effectively doubling or tripling the price paid. The upper-bound shadow prices in our framework, where closed-station consumers face valuations well above the controlled price, correspond to exactly such queuing equilibria.

7.6 Misallocation beyond Geography

A barrel of crude can be refined into a slate of outputs: propane, butane, gasoline, jet fuel, heating oil, petrochemical feedstocks, and others. On the margin, tradeoffs exist: a given quantity of crude oil can yield more gasoline at the expense of less fuel oil, or vice versa. When prices are controlled product-by-product, these tradeoffs interact with small ”nuisance” wedges (refining costs, accounting conventions, transportation frictions, and priority categories) to generate large shifts in the realized fuel mix. During the controls, scarcity therefore appeared not only across space but across products. Propane shortages, diesel shortages, and heating-oil shortages rarely arrived as a single, stationary ”energy shortage”; instead the binding constraint rotated across fuels as demand and costs moved, as the chaos theorem predicts.

A key mechanism was emphasized in 1973 testimony before the Joint Economic Committee[10]:

> Phases 3 and 3A have created a number of price anomalies that frequently prevent the movement of products in proper directions, simply as a result of price considerations.

Heating oil illustrates both product-mix and temporal misallocation. Heating-oil prices were initially frozen at (low-demand) summer 1971 levels, weakening incentives to carry inventories into winter and to direct marginal refinery yield toward distillate. Shortages emerged in the winters of 1972–1973 and 1973–1974, contributing to school closures. Freezing prices led to freezing people (Poole 1973; Verleger 1979; Deacon, Mead, and Agarwal 1980).

Input–output misallocation also arose outside fuels whenever controlled output prices met uncontrolled input costs (or vice versa); Mulligan (2025) formalizes this supply-chain mechanism. In the summer of 1973, chicken farmers gassed, drowned, and suffocated roughly a million baby chicks (TIME 1973; The New York Times 1973; Associated Press 1973). Retail chicken prices were controlled, but feed costs were not. That did more

10. Testimony of James Emison, p. 205. The hearings contain extensive discussion of rotating shortages (U.S. Congress, Joint Economic Committee, Subcommittee on Consumer Economics 1973).

than reduce poultry output: it blocked the price system from allocating supply across production stages. Raising chicks is an intertemporal, input–output segment. Feed today becomes broilers weeks later. With revenues capped along that grow-out path, relatively small increases in feed costs pushed producers to a corner: abandon the “raise chicks into chickens” segment and liquidate immediately. As Madison Clements of Waco put it, “It’s cheaper to drown em than to put em down and raise em," (Associated Press 1973).

The same temporal segmentation appeared in livestock: dairy farmers slaughtered cows and hog farmers culled breeding stock, temporarily increasing meat supply while reducing future flows of milk and pork. In each case, controls flattened relative returns across time and uses, so small cost wedges selected a vertex allocation: all supply routed to the "now" segment and essentially zero to the "later" segment.

Finally, supply-chain complexity amplified inefficiencies in the output mix. Allocation rules designated priority end uses, but policymakers underestimated input-output linkages. Oil production was prioritized, yet some upstream inputs were not. Propane, for instance, was essential for producing plastic piping used in oil extraction, but the plastics industry initially lacked priority designation. The resulting pipe shortages disrupted the very oil production that allocation policy sought to protect (Office 1974).

The misallocation generated by the price controls ran into political economy. Faced with shortage-chaos the political system was pushed toward direct quantity management. By the 1979 episode, for example, the DOE threatened yield regulations requiring refiners to produce a specified share of heating oil per barrel of crude and instructed refiners on inventory accumulation ahead of summer gasoline demand (Bradley 1997; Verleger 1979).

Allocation by fiat is a natural response to shortages because, as the chaos theorem indicates, once prices are prevented from doing their job, “the market” no longer selects a coherent allocation—it delivers corner outcomes with no welfare ordering. Price controls therefore tend to metastasize into quantity controls. But quantity mandates do not escape the chaos theorem: the mandated allocations are themselves fragile, because small errors in the parameters the regulator uses to set them can flip the outcome between corners.[11]

8 Conclusion

In markets without price controls, arbitrage pushes goods toward their highest-value uses, equalizing shadow prices across space, time, and sector. A binding price ceiling breaks this. Once money prices are frozen, the incentive to move goods to higher-value segments vanishes, and the efficient allocation becomes a knife-edge point inside a much larger feasible set.

The Chaos Theorem formalizes the consequence: equilibrium generically lands on corners where some markets are fully served and others receive nothing. Infinitesimal changes in nuisance parameters such as transportation costs can flip the economy between corner allocations, generating discontinuous welfare jumps from continuous parameter shifts. The patchwork of the 1973–74 gasoline crisis is not an accident but the generic outcome when prices cannot signal scarcity.

Our robust bounds analysis quantifies these welfare costs without parametric demand assumptions. Using only observed allocations, the ceiling price, and empirically plausible slope bounds, we find misallocation losses on the order of 1 to 9 times the Harberger

11. Communism can be understood as universal price controls combined with full quantity allocation. A full analysis is beyond the scope of this paper, but the shortage economy that Kornai (1992) documented exhibited dynamics strikingly consistent with the chaos theorem. Soviet citizens carried avoska bags—from avos', meaning "perhaps" or "just in case"—because shortages were unpredictable: shoes today, soap tomorrow, nothing next week. People joined any queue they encountered, often without knowing what was being sold, because the queue itself signaled temporary availability (Smith 1984). Production exhibited parallel irregularities. Factories "stormed"—alternating between slack periods when little was accomplished and frantic bursts of activity as plan deadlines approached (Filtzer 1996)—because input deliveries were erratic and uncoordinated with production schedules. The common mechanism was the absence of price signals to coordinate dispersed decisions, producing systemwide patterns of shortage and surplus that shifted unpredictably in response to small perturbations.

triangle; the quantity reduction accounts for a minority of total welfare loss.

The same logic extends beyond geography. Heating-oil prices frozen at summer levels eliminated incentives to store winter inventories; controlled chicken prices with uncontrolled feed costs drove farmers to a corner on the input–output, intertemporal margin, leading them to destroy over a million chicks. Whenever controls suppress price variation across segments, seller indifference turns small cost wedges into large misallocations. The shortage-chaos leads naturally to quantity controls.

The broader lesson is general: whenever a ceiling fragments an integrated market—gasoline, rental housing, agriculture, medical care—the main cost is not the familiar triangle, but the hidden misallocation behind it.

A Omitted Proofs

A.1 Theorem 1 Proof

Proof of Theorem 1. Because $\widetilde{W}_i(q_i) = \int_0^{q_i} P_i(x)\,dx$ is concave (as an integral of a nonincreasing function), the sum $\widetilde{W} = \sum_i \widetilde{W}_i$ is concave on the convex polytope $\mathcal{F}$. A concave function attains its minimum at an extreme point, proving (i).

For (ii), fix a minimizer $q^w \in \mathcal{F}$ of $\widetilde{W}$. Since each P_i is continuous on $[0, \bar{q}_i]$, $\widetilde{W}$ is continuously differentiable on F, with $\frac{\partial \widetilde{W}}{\partial q_i}(q) = P_i(q_i)$ for all $q \in \mathcal{F}$. The standard Karush-Kuhn-Tucker conditions apply: at q^w, there exist $\lambda, \{\mu_i\}, \{\nu_i\}$ such that

$$
\begin{aligned}
&\text{(stationarity)} && P_i(q_i^w) + \lambda - \mu_i + \nu_i = 0 && \text{for all } i, && (A1)\\
&\text{(complementary slackness)} && \mu_i\, q_i^w = 0, \qquad \nu_i\,(q_i^w - \bar{q}_i) = 0 && \text{for all } i, && (A2)\\
&\text{(primal feasibility)} && \sum_i q_i^w = \bar{Q}, \quad 0 \le q_i^w \le \bar{q}_i && \text{for all } i, && (A3)\\
&\text{(dual feasibility)} && \mu_i, \nu_i \ge 0 && \text{for all } i. && (A4)
\end{aligned}
$$

We define $\lambda_{\text{worst}} := -\lambda$ and observe that there are three possible cases for each market i.

We go through them one-by-one.

Case 1: $0 < q_i^w < \bar{q}_i$**.** Both constraints $-q_i \le 0$ and $q_i - \bar{q}_i \le 0$ are slack at q^w, so by complementary slackness ($A2$) we must have $\mu_i = 0 = \nu_i = 0$. We substitute these into ($A1$) to deduce $P_i(q_i^w) = \lambda_{\text{worst}}$.

Case 2: $q_i^w = 0$**.** Here the constraint $q_i - \bar{q}_i \leq 0$ is slack (since $q_i^w - \bar{q}_i < 0$), so $\nu_i = 0$ by ($A2$), while $\mu_i \geq 0$ is free. Equation ($A1$) becomes $P_i(0) = \lambda_{\text{worst}} + \mu_i \geq \lambda_{\text{worst}}$.

Case 3: $q_i^w = \bar{q}_i$**.** Now the constraint $-q_i \leq 0$ is slack (since $-q_i^w = -\bar{q}_i < 0$), so $\mu_i = 0$ by ($A2$), while $\nu_i \geq 0$ is free. Equation ($A1$) becomes $P_i(\bar{q}_i) = \lambda_{\text{worst}} - \nu_i \leq \lambda_{\text{worst}}$.

As each $P_i(\cdot)$ is weakly decreasing and nonnegative, necessarily $\lambda_{\text{worst}} \geq 0$.

Part (iii) follows immediately: for a single-market allocation with $q_j = \bar{Q}$ and $q_{k \neq j} = 0$, we have $\widetilde{W} = \int_0^{\bar{Q}} P_j(x)dx$, so minimizing across j yields the stated criterion. ■

A.2 Chaos Theorem

This appendix provides the technical machinery for Theorem 2.

Let $\theta \in \Theta \subset \mathbb{R}^m$ be a compact, nonempty parameter space. There are submarkets $i = 1, \ldots, n$ with inverse demands $P_i(x;\theta)$. We assume i) for each i and every θ, the map $x \mapsto P_i(x;\theta)$ is continuous and strictly decreasing; and ii) for each x, the map $\theta \mapsto P_i(x;\theta)$ is continuous. We also assume iii) submarket unit costs $c_i(\theta)$ and the total quantity $\bar{Q}(\theta) > 0$ are continuous in θ.

The feasible splits $q = (q_1, \ldots, q_n)$ satisfy $\sum_{i=1}^n q_i = \bar{Q}(\theta)$ and $q_i \geq 0$. We define

$$W(q;\theta) = \sum_{i=1}^{n} \int_0^{q_i} P_i(x;\theta)dx.$$

Consider

$$\max_{q \geq 0} \sum_{i=1}^{n} \int_0^{q_i} P_i(x;\theta)dx - \sum_{i=1}^{n} c_i(\theta)q_i \quad \text{s.t.} \quad \sum_{i=1}^{n} q_i = \bar{Q}(\theta).$$

As each $P_i(\cdot;\theta)$ is strictly decreasing, the objective is strictly concave in q, hence the maximizer $q^*(\theta)$ is unique. The feasible correspondence

$$\Gamma(\theta) := \left\{ q \in \mathbb{R}_+^n : \sum_i q_i = \bar{Q}(\theta) \right\},$$

has compact values and a closed graph by the continuity of $\bar{Q}$. The objective is continuous in (q,θ) (by the joint continuity of P_i and continuity of c_i and $\bar{Q}$). Therefore, by Berge's

Maximum Theorem, both the unique maximizer $\theta \mapsto q^*(\theta)$ and the value function

$$\theta \mapsto \sum_{i=1}^{n} \int_0^{q_i^*(\theta)} P_i(x;\theta)dx - \sum_{i=1}^{n} c_i(\theta) q_i^*(\theta),$$

are continuous in θ. On any region where the optimizer is interior, the multiplier $\theta \mapsto \lambda(\theta)$ is also continuous, as $\lambda(\theta) = P_i(q_i^*(\theta);\theta) - c_i(\theta)$.

Now suppose we have a binding ceiling $\bar{p}$. Define $\bar{q}_i(\theta) := D_i(\bar{p};\theta)$ and aggregate quantity $\bar{Q}(\theta) := S(\bar{p};\theta)$. The feasible set is

$$\mathcal{F}(\theta) = \left\{ q \in \mathbb{R}_+^n : \sum_{i=1}^{n} q_i = \bar{Q}(\theta), \, 0 \le q_i \le \bar{q}_i(\theta) \; \forall i \right\}.$$

We make the genericity assumptions

I. $0 < \bar{Q}(\theta) < \sum_{i=1}^{n} \bar{q}_i(\theta)$;

II. $\bar{Q}(\theta) \neq \sum_{i \in S} \bar{q}_i(\theta)$ for every $S \subset \{1, \ldots, n\}$;

III. $c_i(\theta) \neq c_j(\theta)$ for $i \neq j$.

We solve

$$\min_{q \in \mathcal{F}(\theta)} c(\theta) \cdot q.$$

We order costs as $c_{(1)}(\theta) < \cdots < c_{(n)}(\theta)$ and let k denote the unique index with $\sum_{s=1}^{k-1} \bar{q}_{(s)}(\theta) < \bar{Q}(\theta) < \sum_{s=1}^{k} \bar{q}_{(s)}(\theta)$. Then the unique optimizer is

$$q_{(r)}^*(\theta) = \begin{cases} \bar{q}_{(r)}(\theta), & r < k, \\ \bar{Q}(\theta) - \sum_{s=1}^{k-1} \bar{q}_{(s)}(\theta), & r = k, \\ 0, & r > k, \end{cases}$$

i.e., we "fill in increasing c" with at most one partial coordinate. Changing θ that breaks a cost tie ($c_i(\theta) = c_j(\theta)$) that changes which markets are filled causes a discrete jump in allocations.

We suppress dependence on the parameter θ and recall the optimization problem, linear program

$$\min \{ c \cdot q : \; q \in \mathcal{F} \} . \tag{$\star$}$$

For the local instability argument we fix the feasible polytope $\mathcal{F} = \mathcal{F}(\theta^*)$ and consider perturbations that act only through the cost vector $c(\theta)$. Thus, we treat c as the effective parameter in ($\star$). We maintain Assumptions I. and II. Assumption II. means that every vertex v of $\mathcal{F}$ has exactly one coordinate strictly between its bounds (the partial coordinate), and all other coordinates are at 0 or at $\bar{q}_i$. To elaborate, a point $v \in \mathcal{F}$ is a vertex iff exactly $n-1$ $q_i \in \{0, \bar{q}_i\}$, with the equality $\sum q_i = \bar{Q}$ then fixing the partial coordinate.

We say that two distinct vertices v and w are *adjacent* if they share an edge: along the edge exactly the pair (r, s) of coordinates moves (with all others fixed at bounds) and

$$wb - v = \Delta(e_s - e_r) \qquad \text{for some } r \neq s \text{ and } \Delta \neq 0, \tag{1}$$

where r is the unique partial index at v (so $0 < v_r < \bar{q}_r$); and at w, one of r or s is the new partial (non-bound coordinate) and the other hits a bound. Equivalently, edges are exactly the one-dimensional faces on which two coordinates are free (subject to the sum constraint), and the others are fixed at bounds.

For any $v \in \mathcal{F}$, the (outer) normal cone is

$$\mathcal{N}_v = \{c \in \mathbb{R}^n : \ c \cdot (x - v) \geq 0 \ \forall x \in F\}.$$

The following is standard:

Remark 4. v solves ($\star$) if and only if $c \in \mathcal{N}_v$.

Moreover,

Remark 5. If a given c makes exactly two distinct vertices v, w optimal, then v, w are adjacent and $c \cdot (w - v) = 0$.

Denote the set of such common-optimal costs by

$$\mathcal{H}_{vw} := \mathcal{N}_v \cap \mathcal{N}_w \cap \{c : \ c \cdot (w - v) = 0\}.$$

We say that $c^\star$ is a *simple tie* for v, w if

$$c^\star \in \operatorname{relint}(\mathcal{H}_{vw}) \quad \text{and} \quad c^\star \notin \mathcal{N}_u \text{ for any other vertex } u.$$

Note that if c makes v and w optimal but the tie is not simple, then every neighborhood of c contains some $\hat{c}$ that is a simple tie for v, w; thus Lemma 1 still yields crossings arbitrarily close to any (v, w) tie.

Lemma 1. *Let $v \neq w$ be adjacent vertices and $c^\star$ a simple tie for v, w. Then for any neighborhood $U \subset \mathbb{R}^n$ of $c^\star$ there exists $\varepsilon > 0$ and a continuously-differentiable function $c\colon [-\varepsilon, \varepsilon] \to U$ such that the unique optimizer of ($\star$) equals v for $t < 0$ and w for $t > 0$, while at $t = 0$ both v and w are optimal.*

Proof. As $c^\star \in \operatorname{relint} \mathcal{H}_{vw}$ and no third cone meets at $c^\star$, there exists $\rho > 0$ with $B(c^\star, \rho) \cap \operatorname{fan}(\mathcal{F}) = \operatorname{int} \mathcal{N}_v \dot\cup \mathcal{H}_{vw} \dot\cup \operatorname{int} \mathcal{N}_w$ ($\dot\cup$ denotes the disjoint union). Thus, for any η with $(w - v) \cdot \eta < 0$, the line $c(t) = c^\star + t\eta$ (for $|t|$ small) satisfies $(w - v) \cdot c(t) = t(w - v) \cdot \eta$ and thus crosses $\mathcal{H}_{vw}$ with a sign change, so $c(t) \in \operatorname{int} \mathcal{N}_v$ for $t < 0$, $c(0) \in \mathcal{H}_{vw}$, and $c(t) \in \operatorname{int} \mathcal{N}_w$ for $t > 0$. Optimality follows from Remark 4. ∎

We continue to let v, w be adjacent vertices with $w - v = \Delta(e_s - e_r)$, $\Delta \neq 0$, as in (1). Let $c(t)$ be a continuously-differentiable function as in Lemma 1. Define the welfare jump at the crossing by

$$\Delta W = \lim_{t \downarrow 0} W(w; c(t)) - \lim_{t \uparrow 0} W(v; c(t)).$$

The next lemma reassures us that the $-c \cdot q$ contribution cancels in the limit, as $c(\cdot)$ is continuous at 0 and $c(0) \cdot (w - v) = 0$.

Lemma 2. *We have*

$$\Delta W = \int_{v_s}^{w_s} P_s(x) dx - \int_{w_r}^{v_r} P_r(x) dx. \tag{2}$$

Proof. By definition,

$$\begin{aligned}
\Delta W &= \lim_{t \downarrow 0} \left(\sum_i \int_0^{w_i} P_i(x) dx - c(t) \cdot w \right) - \lim_{t \uparrow 0} \left(\sum_i \int_0^{v_i} P_i(x) dx - c(t) \cdot v \right) \\
&= \sum_i \left(\int_0^{w_i} P_i - \int_0^{v_i} P_i \right) - \lim_{t \to 0} c(t) \cdot (w - v) \\
&= \sum_i \int_{v_i}^{w_i} P_i - c(0) \cdot (w - v) = \int_{v_s}^{w_s} P_s - \int_{w_r}^{v_r} P_r,
\end{aligned}$$

where we used the definition of a simple tie and Remark 5 to eliminate $c(0) \cdot (w - v) = 0$. ∎

Proof of Theorem 2

Proof of Theorem 2. By Lemma 1, we can construct a smooth crossing $c(\cdot)$. By Lemma 2, $\Delta W[c] = \int_{v_s}^{w_s} P_s - \int_{w_r}^{v_r} P_r$, which is nonzero by hypothesis ($\Delta W \neq 0$).

Define the reversed crossing $\tilde{c}(t) := c(-t)$. The optimizer flips in the opposite direction, and $\Delta W[\tilde{c}] = -\Delta W[c]$, with the same magnitude $|\Delta W|$.

Setting $c_+ := c$ and $c_- := \tilde{c}$ yields the two paths with welfare jumps $+|\Delta W|$ and $-|\Delta W|$ respectively. Since these have opposite signs, welfare is not locally monotone in any neighborhood of $c^* \equiv c(\theta^*)$. ■

Proof of Corollary 1. For every neighborhood U of $c^* \equiv c(\theta^*)$, Lemma 1 provides $c^-, c^+ \in U$ with $q^*(c^-) = v$ and $q^*(c^+) = w$. Then

$$\|q^*(c^+) - q^*(c^-)\|_1 = \|w - v\|_1 = |\Delta| + |-\Delta| = 2\Delta > 0. \qquad \blacksquare$$

B Proposition 3 and Theorem 7 Formal Statements and Proofs

Recall our assumption:

Assumption 2. For all i, P_i is nonincreasing and satisfies

I. $P_i(q_i^{\text{obs}}) = p_{0,i} = \bar{p} - b_i \quad \text{with} \quad b_i \in [\underline{b}_i, \bar{b}_i]$.
II. There exist $g_{i,L} < g_{i,U} < 0$ with $g_{i,L} \leq P_i'(q) \leq g_{i,U}$ for a.e. $q \in [0, q_i^{\max}]$.
III. $P_i(0) \leq M_i < \infty$
IV. $P_i \in W^{1,\infty}([0, q_i^{\max}])$, the space of bounded Lipschitz-continuous functions on $[0, q_i^{\max}]$.

Recall also that $q_i(\cdot)$ denotes the left-continuous generalized inverse of $P_i(\cdot)$:

$$q_i(p) := \begin{cases} \inf\{x \in [0, q_i^{\max}] : P_i(x) \leq p\}, & \text{if the set is nonempty,} \\ q_i^{\max}, & \text{otherwise.} \end{cases}$$

We first compute easy bounds for q_i. By the fundamental theorem of calculus, $P_i(q) = p_{0,i} + \int_{q_i^{\text{obs}}}^{q} P_i'(s)ds$. By Assumption 2II., $P_i(q)$ lies between the two affine functions $p_{0,i} + g_{i,L}(q - q_i^{\text{obs}})$ and $p_{0,i} + g_{i,U}(q - q_i^{\text{obs}})$. Rearranging these inequalities yields

$$\min\left\{q_i^{\text{obs}} + \frac{p - p_{0,i}}{g_{i,U}},\ q_i^{\text{obs}} + \frac{p - p_{0,i}}{g_{i,L}}\right\} \leq q_i(p) \leq \max\left\{q_i^{\text{obs}} + \frac{p - p_{0,i}}{g_{i,U}},\ q_i^{\text{obs}} + \frac{p - p_{0,i}}{g_{i,L}}\right\}.$$

To incorporate the choke-price restriction $P_i(0) \leq M_i$ (Assumption 2III.), we explicitly enforce that demand is zero at prices $p \geq M_i$ and rule out (q, p) pairs that cannot be connected to $q = 0$ without exceeding M_i under the least-negative slope $g_{i,U}$.

We summarize the bounds with two functions. For $p \in \mathbb{R}$, we define

$$\ell_i(p) := \begin{cases} 0, & p \geq M_i, \\ \max\left\{0,\ \min\left\{q_i^{\text{obs}} + \frac{p-p_{0,i}}{g_{i,U}},\ q_i^{\text{obs}} + \frac{p-p_{0,i}}{g_{i,L}}\right\}\right\}, & p < M_i, \end{cases}$$

and

$$u_i(p) := \begin{cases} 0, & p \geq M_i, \\ \min\left\{q_i^{\max},\ \max\left\{q_i^{\text{obs}} + \frac{p-p_{0,i}}{g_{i,U}},\ q_i^{\text{obs}} + \frac{p-p_{0,i}}{g_{i,L}}\right\},\ \frac{M_i-p}{-g_{i,U}}\right\}, & p < M_i. \end{cases}$$

We also characterize set of candidate common shadow prices consistent with the aggregate quantity constraint as follows. Define $L(p) := \sum_i \ell_i(p)$, $U(p) := \sum_i u_i(p)$, and

$$\mathcal{I} := \left\{p \in \mathbb{R} \colon L(p) \leq \bar{Q} \leq U(p)\right\}.$$

Then $q_i(p) \in [\ell_i(p), u_i(p)]$ for all $p \in \mathcal{I}$.

To summarize things: at price p, $\ell_i(p)$ and $u_i(p)$ are the smallest and largest (respectively) quantities in market i that are consistent with $q_i(\cdot)$ (i) passing through $\left(q_i^{\text{obs}}, p_{0,i}\right)$, (ii) obeying the slope bounds, and (iii) satisfying the choke-price restriction $P_i(0) \leq M_i$. The set $\mathcal{I}$ is the set of candidate common shadow prices p for which the interval constraints across markets are compatible with $\sum_i q_i(p) = \bar{Q}$.

Next, for $P \in \mathcal{P}$, define

$$\Phi(P) := \sum_{i=1}^{n} \int_{q_i^{\text{obs}}}^{q_i^*(P)} \left(P_i(x) - p^*(P)\right) dx,$$

where $q^*(P)$ is the equal-shadow-price split of total quantity $\bar{Q}$ and $p^*(P)$ is its Lagrange multiplier. Note that $\sum_i q_i^*(P) = \bar{Q}$ and $\sum_i q_i^{\text{obs}} = \bar{Q}$, so using $p^*(P)$ (rather than $\bar{p}$) is harmless since $\sum_i \left(q_i^*(P) - q_i^{\text{obs}}\right) = 0$. Recall $\overline{\Phi} := \max_{P \in \mathcal{P}} \Phi(P)$, and $\underline{\Phi} := \min_{P \in \mathcal{P}} \Phi(P)$.

We make the following interiority assumption

Assumption 3. The set $\mathcal{I}$ is nonempty and compact. Moreover, for every i and every

$p \in \mathcal{I}$, and every $s \in [\min\{p, p_{0,i}\}, \max\{p, p_{0,i}\}]$,

$$0 < q_i^{\text{obs}} + \frac{s - p_{0,i}}{g_{i,U}} < q_i^{\max}, \qquad 0 < q_i^{\text{obs}} + \frac{s - p_{0,i}}{g_{i,L}} < q_i^{\max}, \qquad s < M_i,$$

and

$$\max\left\{q_i^{\text{obs}} + \frac{s - p_{0,i}}{g_{i,U}},\ q_i^{\text{obs}} + \frac{s - p_{0,i}}{g_{i,L}}\right\} < \frac{M_i - s}{-g_{i,U}}.$$

For our purposes, this assumption means that $\ell_i(\cdot)$ and $u_i(\cdot)$ simplify to

$$\ell_i(p) = \min\left\{q_i^{\text{obs}} + \frac{p - p_{0,i}}{g_{i,U}},\ q_i^{\text{obs}} + \frac{p - p_{0,i}}{g_{i,L}}\right\}, \tag{A1}$$

and

$$u_i(p) = \max\left\{q_i^{\text{obs}} + \frac{p - p_{0,i}}{g_{i,U}},\ q_i^{\text{obs}} + \frac{p - p_{0,i}}{g_{i,L}}\right\}, \tag{A2}$$

and take values in $\left(0, q_i^{\max}\right)$. Henceforth, we take these simplified expressions to be the definitions of ℓ_i and u_i. Moreover, on $p \in \mathcal{I}$, $q_i(p)$ simplifies to $q_i(p) = P_i^{-1}(p)$.

Next, we introduce several more objects: for each i, define

$$\alpha_i := \frac{1}{g_{i,U}}, \qquad \beta_i := \frac{1}{g_{i,L}}, \quad \text{and} \quad d_i := \beta_i - \alpha_i = \frac{1}{g_{i,L}} - \frac{1}{g_{i,U}} > 0.$$

For each $p \in \mathcal{I}$, define the baseline endpoint choices

$$q_i^+(p) := \begin{cases} \ell_i(p), & p \geq p_{0,i}, \\ u_i(p), & p < p_{0,i}, \end{cases} \quad \text{and} \quad q_i^-(p) := \begin{cases} u_i(p), & p \geq p_{0,i}, \\ \ell_i(p), & p < p_{0,i}. \end{cases}$$

Let

$$\Delta^+(p) := \bar{Q} - \sum_{i=1}^{n} q_i^+(p), \quad \text{and} \quad \Delta^-(p) := \bar{Q} - \sum_{i=1}^{n} q_i^-(p).$$

and define active sets

$$\mathcal{A}^+(p) := \begin{cases} \{i \colon p \geq p_{0,i}\}, & \Delta^+(p) > 0, \\ \{i \colon p < p_{0,i}\}, & \Delta^+(p) < 0, \\ \emptyset, & \Delta^+(p) = 0, \end{cases} \quad \text{and} \quad \mathcal{A}^-(p) := \begin{cases} \{i \colon p < p_{0,i}\}, & \Delta^-(p) > 0, \\ \{i \colon p \geq p_{0,i}\}, & \Delta^-(p) < 0, \\ \emptyset, & \Delta^-(p) = 0. \end{cases}$$

Define also the triangle-penalty value functions

$$\Psi^+(p) := \min_{\{\delta_i\}_{i\in\mathcal{A}^+(p)}} \left\{ \sum_{i\in\mathcal{A}^+(p)} \frac{\delta_i^2}{2d_i} \right\} \quad \text{s.t.} \quad \sum_{i\in\mathcal{A}^+(p)} \delta_i = \left|\Delta^+(p)\right|, \qquad 0 \le \delta_i \le u_i(p) - \ell_i(p),$$

and

$$\Psi^-(p) := \min_{\{\delta_i\}_{i\in\mathcal{A}^-(p)}} \left\{ \sum_{i\in\mathcal{A}^-(p)} \frac{\delta_i^2}{2d_i} \right\} \quad \text{s.t.} \quad \sum_{i\in\mathcal{A}^-(p)} \delta_i = \left|\Delta^-(p)\right|, \qquad 0 \le \delta_i \le u_i(p) - \ell_i(p).$$

For a fixed $p \in \mathcal{I}$ we term the solution to $\max_{\substack{P\in\mathcal{P} \\ p^*(P)=p}} \Phi(P)$ (in turn, $\min_{\substack{P\in\mathcal{P} \\ p^*(P)=p}} \Phi(P)$) an *inner optimizer*.

Proposition 3. *Posit Assumptions 2 and 3. Then*

$$\overline{\Phi} = \max_{p\in\mathcal{I}} \left\{ \bar{Q}p - \sum_{i=1}^{n} q_i^{\text{obs}} p_{0,i} - \sum_{i:\ p\ge p_{0,i}} \int_{p_{0,i}}^{p} \ell_i(s)ds + \sum_{i:\ p<p_{0,i}} \int_{p}^{p_{0,i}} u_i(s)ds - \Psi^+(p) \right\}, \quad (1)$$

and

$$\underline{\Phi} = \min_{p\in\mathcal{I}} \left\{ \bar{Q}p - \sum_{i=1}^{n} q_i^{\text{obs}} p_{0,i} - \sum_{i:\ p\ge p_{0,i}} \int_{p_{0,i}}^{p} u_i(s)ds + \sum_{i:\ p<p_{0,i}} \int_{p}^{p_{0,i}} \ell_i(s)ds + \Psi^-(p) \right\}. \quad (2)$$

Moreover, for any $p \in \mathcal{I}$, the inner optimizer may be chosen so that each inverse-quantity function $q_i(\cdot)$ is continuous and piecewise affine on $[\min\{p, p_{0,i}\}, \max\{p, p_{0,i}\}]$, has a.e. slope in $\{\alpha_i, \beta_i\}$, and is piecewise affine with at most one kink.

We first assemble several lemmas, deferring the proofs of these lemmas to the supplementary appendix.

Lemma 3. *Under Assumptions 2 and 3, the functions $\Psi^+(\cdot)$ and $\Psi^-(\cdot)$ are continuous on $\mathcal{I}$.*

Next, for any $p \in \mathcal{I}$, define $\mathcal{I}_+(p) := \{i\colon p \ge p_{0,i}\}$ and $\mathcal{I}_-(p) := \{i\colon p < p_{0,i}\}$.

Lemma 4. *Posit Assumptions 2 and 3 and fix $p \in \mathcal{I}$. Then, the inner problem of maximizing $\Phi(P)$ over $P \in \mathcal{P}$ satisfying $p^*(P) = p$ is equivalent to minimizing $\int_{p_{0,i}}^{p} q_i$ for $i \in \mathcal{I}_+(p)$ and maximizing $\int_{p}^{p_{0,i}} q_i$ for $i \in \mathcal{I}_-(p)$, subject to feasibility and $\sum_i q_i(p) = \bar{Q}$.*

Lemma 5. *$q_i'(s) \in [\alpha_i, \beta_i]$ for a.e. $s \in [\min\{p, p_{0,i}\}, \max\{p, p_{0,i}\}]$.*

Lemma 6. *Now let $t_0 < t_1$, $\alpha < \beta$, and let function q be absolutely continuous on $[t_0, t_1]$ with $q(t_0) = q_0$ and $q'(t) \in [\alpha, \beta]$ a.e. Define $\underline{q}(t) := q_0 + \alpha(t - t_0)$. If $q(t_1) = \underline{q}(t_1) + \delta$ for some $\delta \in [0, (\beta - \alpha)(t_1 - t_0)]$, then*

$$\int_{t_0}^{t_1} q(t)dt \geq \int_{t_0}^{t_1} \underline{q}(t)dt + \frac{\delta^2}{2(\beta - \alpha)}. \tag{A6}$$

Moreover, equality holds if and only if q is "bang-bang;" $q'(t) = \alpha$ a.e. on $[t_0, t_1 - h]$ and $q'(t) = \beta$ a.e. on $[t_1 - h, t_1]$, where $h = \delta / (\beta - \alpha)$.

Proof of Proposition 3. Fix $p \in \mathcal{I}$ and consider the inner maximization problem of maximizing $\Phi(P)$, which by Lemma 4 is equivalent to minimizing $\int_{p_{0,i}}^{p} q_i$ for $i \in \mathcal{I}_+(p)$ and maximizing $\int_{p}^{p_{0,i}} q_i$ for $i \in \mathcal{I}_-(p)$, subject to $\sum_i q_i(p) = \bar{Q}$. Ignoring the constraint $\sum_i q_i(p) = \bar{Q}$ and optimizing the integral term in (A5) with $p^* = p$, $\sum_{i=1}^{n} \int_{p_{0,i}}^{p^*} q_i(s)ds$, the solution is $q_i = \ell_i$ for $i \in \mathcal{I}_+(p)$ and $q_i = u_i$ for $i \in \mathcal{I}_-(p)$, corresponding to $q_i^+(p)$.

If this sum does not equal $\bar{Q}$, the constraint $\sum_i q_i(p) = \bar{Q}$ forces deviations from the endpoints (ℓ_i and u_i) at p on exactly the set $\mathcal{A}^+(p)$, as follows (with Claim 6 formally proved in the supplementary appendix). Note that $\sum_i q_i^+(p) \neq \bar{Q}$ if and only if $\Delta^+(p) \neq 0$.

Claim 6. *Fix $p \in \mathcal{I}$.*

I. *If $\Delta^+(p) > 0$, there exist numbers $\{\delta_i\}_{i \in \mathcal{I}_+(p)}$ such that $0 \leq \delta_i \leq u_i(p) - \ell_i(p)$ for all $i \in \mathcal{I}_+(p)$ and $\sum_{i \in \mathcal{I}_+(p)} \delta_i = \Delta^+(p)$, and the endpoint vector $q(p)$ defined by*

$$q_i(p) = \ell_i(p) + \delta_i \text{ for } i \in \mathcal{I}_+(p), \qquad q_i(p) = u_i(p) \text{ for } i \in \mathcal{I}_-(p)$$

satisfies $\sum_{i=1}^{n} q_i(p) = \bar{Q}$. Moreover, in the inner problem, there is an optimizer with $q_i(p) = u_i(p)$ for all $i \in \mathcal{I}_-(p)$; equivalently, deviations from the baseline endpoints $q_i^+(p)$ may be taken without loss of optimality to occur only on $\mathcal{A}^+(p) = \mathcal{I}_+(p)$.

II. *If $\Delta^+(p) < 0$, there exist numbers $\{\delta_i\}_{i \in \mathcal{I}_-(p)}$ such that $0 \leq \delta_i \leq u_i(p) - \ell_i(p)$ for all $i \in \mathcal{I}_-(p)$ and $\sum_{i \in \mathcal{I}_-(p)} \delta_i = -\Delta^+(p)$, and the endpoint vector $q(p)$ defined by*

$$q_i(p) = \ell_i(p) \text{ for } i \in \mathcal{I}_+(p), \qquad q_i(p) = u_i(p) - \delta_i \text{ for } i \in \mathcal{I}_-(p)$$

satisfies $\sum_{i=1}^{n} q_i(p) = \bar{Q}$. Moreover, in the inner problem, there is an optimizer with $q_i(p) = \ell_i(p)$ for all $i \in \mathcal{I}_+(p)$; equivalently, deviations from the baseline endpoints

$q_i^+(p)$ may be taken without loss of optimality to occur only on $\mathcal{A}^+(p) = \mathcal{I}_-(p)$.

Now fix $p \in \mathcal{I}$. By Claim 6, we may restrict attention to endpoint deviations supported on $\mathcal{A}^+(p)$. Let $\{\delta_i\}_{i\in\mathcal{A}^+(p)}$ be feasible for $\Psi^+(p)$, and set $\delta_i := 0$ for $i \notin \mathcal{A}^+(p)$. Then $q_i(p) = \ell_i(p) + \delta_i$ for all i with $p \geq p_{0,i}$, and $q_i(p) = u_i(p) - \delta_i$ for all i with $p < p_{0,i}$.

By Lemma 5, $q_i'(s) \in [\alpha_i, \beta_i]$ a.e. on $s \in [\min\{p, p_{0,i}\}, \max\{p, p_{0,i}\}]$, so if $p \geq p_{0,i}$ and $q_i(p) = \ell_i(p) + \delta_i$, we can apply Lemma 6 with $\alpha = \alpha_i$, $\beta = \beta_i$, $t_0 = p_{0,i}$, $t_1 = p$ to obtain

$$\int_{p_{0,i}}^{p} q_i(s)ds \geq \int_{p_{0,i}}^{p} \ell_i(s)ds + \frac{\delta_i^2}{2d_i}. \tag{$A8$}$$

Analogously, if $p < p_{0,i}$ and $q_i(p) = u_i(p) - \delta_i$, then by a change of variables reversing the integration limits, applying Lemma 6 yields

$$\int_{p}^{p_{0,i}} q_i(s)ds \leq \int_{p}^{p_{0,i}} u_i(s)ds - \frac{\delta_i^2}{2d_i}. \tag{A9}$$

Using Lemma 4 and substituting ($A8$) and (A9) into (A5) shows that for fixed $p \in \mathcal{I}$,

$$\sup_{\substack{P\in\mathcal{P}\\ p^*(P)=p}} \Phi(P) \leq \bar{Q}p - \sum_{i=1}^{n} q_i^{\text{obs}} p_{0,i} - \sum_{i:\, p\geq p_{0,i}} \int_{p_{0,i}}^{p} \ell_i(s)ds + \sum_{i:\, p<p_{0,i}} \int_{p}^{p_{0,i}} u_i(s)ds \\ - \Psi^+(p). \tag{$A10$}$$

Continue to fix $p \in \mathcal{I}$ and let $\{\delta_i^*\}_{i\in\mathcal{A}^+(p)}$ attain $\Psi^+(p)$ (existence follows from compactness of the feasible set and continuity of the objective). Extend by setting $\delta_i^* := 0$ for $i \notin \mathcal{A}^+(p)$, and define $h_i := \delta_i^*/d_i$ for each i. Construct $q_i^*(\cdot)$ on $[\min\{p, p_{0,i}\}, \max\{p, p_{0,i}\}]$ by the equality cases in ($A8$) and (A9):

- If $p \geq p_{0,i}$, set $q_i^*(s) = \ell_i(s)$ for $s \in [p_{0,i}, p - h_i]$, and set

$$q_i^*(s) := \ell_i(p - h_i) + \beta_i(s - (p - h_i)) \quad \text{for } s \in [p - h_i, p].$$

- If $p < p_{0,i}$, set $q_i^*(s) = u_i(s)$ for $s \in [p + h_i, p_{0,i}]$, and set

$$q_i^*(s) := u_i(p + h_i) + \beta_i(s - (p + h_i)) \quad \text{for } s \in [p, p + h_i].$$

By construction, each $q_i^*(\cdot)$ is continuous and piecewise affine, has a.e. slope in $\{\alpha_i, \beta_i\}$, and satisfies $q_i^*(p_{0,i}) = q_i^{\text{obs}}$. Moreover, $\sum_{i=1}^{n} q_i^*(p) = \sum_{i=1}^{n} q_i^+(p) + \Delta^+(p) = \bar{Q}$, and q^*

attains equality in ($A8$) and (A9). Therefore the upper bound ($A10$) is tight and $q^*(\cdot)$ attains the inner optimum at this p.

Define P_i^* as the inverse of q_i^*. Then P_i^* is continuous and piecewise affine on $\left[0, q_i^{\max}\right]$, with a.e. derivative in $\{g_{i,L}, g_{i,U}\}$, satisfies $P_i^*\left(q_i^{\text{obs}}\right) = p_{0,i}$, and satisfies $P_i^*(0) \leq M_i$ because $q_i^*(p) = 0$ for all $p \geq M_i$ is feasible under the choke-aware bounds. Thus, $P^* \in \mathcal{P}$ attains the inner optimum at p.

Finally, we maximize over $p \in \mathcal{I}$. Define $F^+(p)$ as the bracketed expression in (1). By Lemma 3, $\Psi^+(p)$ is continuous in p, so $F^+(p)$ is continuous on the compact set $\mathcal{I}$ and attains a maximum, proving (1).

The argument for the lower bound $\underline{\Phi}$ is identical, *mutatis mutandis*. ■

Theorem 7. *Posit Assumptions 2 and 3. Both $\overline{\Phi}$ and $\underline{\Phi}$ are attained by some $P^* \in \mathcal{P}$. Moreover, each P_i^* may be chosen continuous and piecewise affine on $\left[0, q_i^{\max}\right]$, with a.e. derivative $P_i^{*\prime}(q) \in \{g_{i,L}, g_{i,U}\}$, and satisfying $P_i^*(0) \leq M_i$.*

Proof. By Proposition 3, the objective in (1) is continuous on the compact set $\mathcal{I}$, so it admits a maximizer $p^* \in \mathcal{I}$. Let $\left\{\delta_i^*\right\}$ attain $\Psi^+(p^*)$. The proof of Proposition 3 constructs inverse-quantity functions $q_i^*(\cdot)$ that attain the inner optimum at p^* and are continuous and piecewise affine with slopes in $\{\alpha_i, \beta_i\}$ on $[\min\{p^*, p_{0,i}\}, \max\{p^*, p_{0,i}\}]$, with at most one slope switch. Outside this interval, extend each $q_i^*(\cdot)$ to all prices by any monotone continuous continuation that keeps slopes in $\{\alpha_i, \beta_i\}$ while $q_i^*(p) > 0$, and then sets $q_i^*(p) = 0$ for all $p \geq M_i$. This extension does not change Φ, since Φ depends on P_i only on $\left[\min\left\{q_i^{\text{obs}}, q_i^*(p^*)\right\}, \max\left\{q_i^{\text{obs}}, q_i^*(p^*)\right\}\right]$. Define P_i^* as the generalized inverse of this extended q_i^*. Then P_i^* is continuous and piecewise affine on $\left[0, q_i^{\max}\right]$, with a.e. derivative in $\{g_{i,L}, g_{i,U}\}$, satisfies $P_i^*\left(q_i^{\text{obs}}\right) = p_{0,i}$, and satisfies $P_i^*(0) \leq M_i$. Therefore $P^* \in \mathcal{P}$ attains $\overline{\Phi}$. The construction for $\underline{\Phi}$ is analogous using (2) and $\Psi^-(p)$. ■

C Additional Robustness Figures

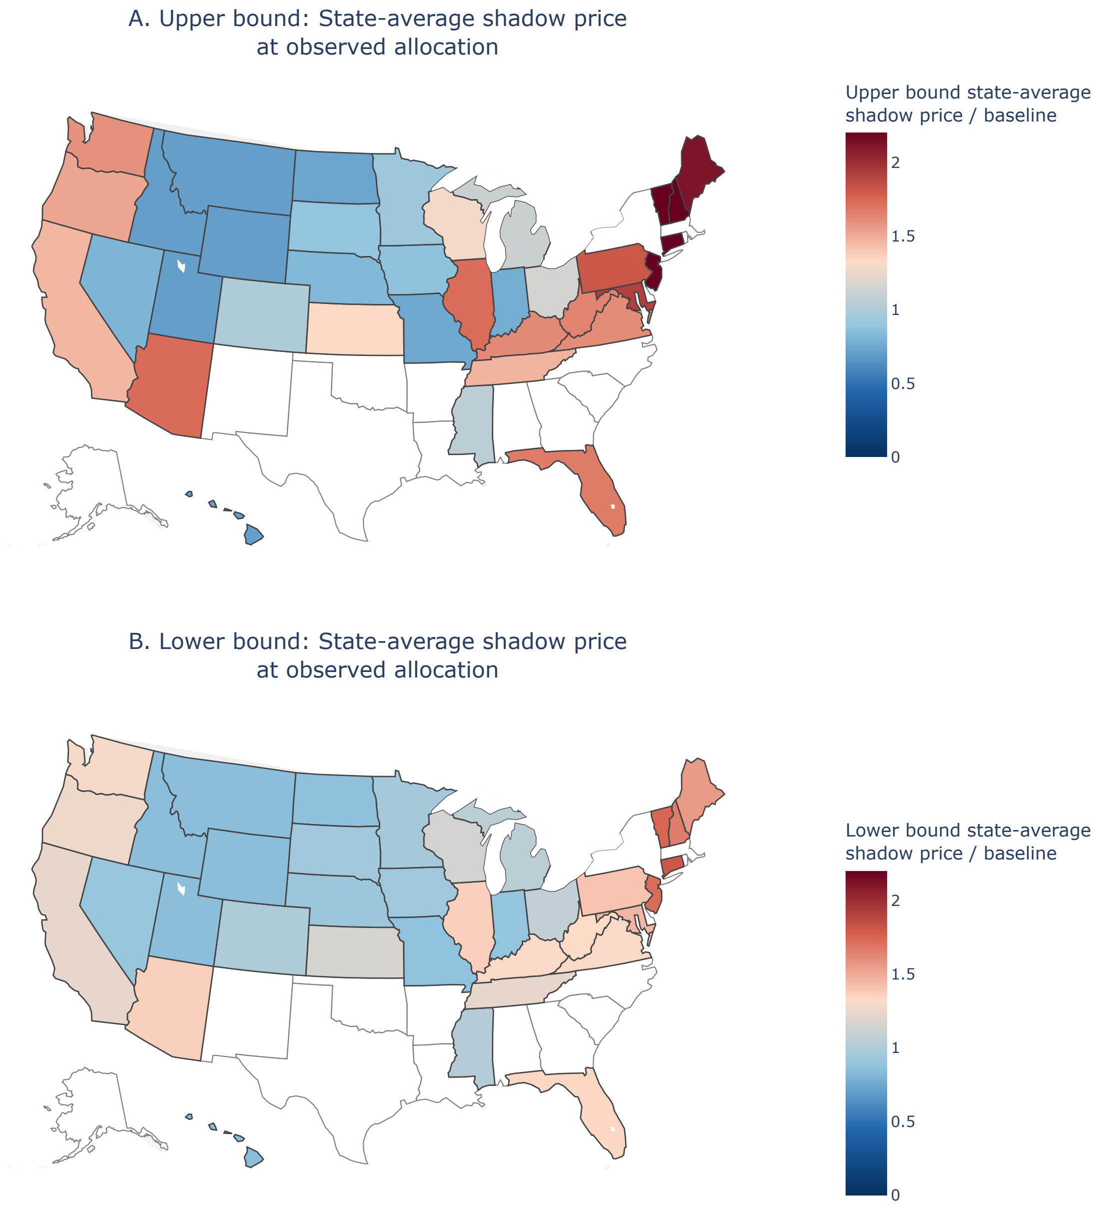

Figure 8: **State-Average Shadow Prices: No Imputation.** Same construction as Figure 7, but using only the 36 states with observed 1972 gallon sales (no imputation). States shown in white lack gallon data. The pattern is similar to the main specification with imputed states included.

D Supplementary Appendix

D.1 Proofs of Lemmas

Lemma 3. *Under Assumptions 2 and 3, the functions $\Psi^+(\cdot)$ and $\Psi^-(\cdot)$ are continuous on $\mathcal{I}$.*

Proof of Lemma 3. We prove continuity of Ψ^+; the argument for Ψ^- is identical.

Under Assumption 3, on $\mathcal{I}$, $\ell_i(\cdot)$ and $u_i(\cdot)$ are given by ($A1$) and ($A2$), and so are continuous on $\mathcal{I}$. Evidently, so too are $q_i^+(\cdot)$, $q_i^-(\cdot)$, $\Delta^+(\cdot)$ and $\Delta^-(\cdot)$.

For each i and $p \in \mathcal{I}$, define the continuous functions

$$\underline{z}_i^+(p) := \ell_i(p) - q_i^+(p), \qquad \text{and} \qquad \overline{z}_i^+(p) := u_i(p) - q_i^+(p).$$

Note that $\underline{z}_i^+(p) \le 0 \le \overline{z}_i^+(p)$ for all i, and

$$\sum_{i=1}^n \underline{z}_i^+(p) = L(p) - \sum_{i=1}^n q_i^+(p) \le \bar{Q} - \sum_{i=1}^n q_i^+(p) = \Delta^+(p) \le U(p) - \sum_{i=1}^n q_i^+(p) = \sum_{i=1}^n \overline{z}_i^+(p),$$

where the inequality uses $p \in \mathcal{I}$, i.e., $L(p) \le \bar{Q} \le U(p)$. Thus, the set

$$Z^+(p) := \left\{ z \in \mathbb{R}^n : \sum_{i=1}^n z_i = \Delta^+(p),\ \underline{z}_i^+(p) \le z_i \le \overline{z}_i^+(p)\ \forall i \right\}$$

is nonempty and compact. Define the optimization-problem

$$\widetilde{\Psi}^+(p) := \min_{z \in Z^+(p)} \sum_{i=1}^n \frac{z_i^2}{2d_i}. \tag{$A3$}$$

Claim 8. $\widetilde{\Psi}^+(p) = \Psi^+(p)$.

Proof. If $\Delta^+(p) = 0$, both values equal 0. If $\Delta^+(p) > 0$, let z be a minimizer of $\widetilde{\Psi}^+(p)$. If some coordinate satisfies $z_i < 0$, define $N := \{i : z_i < 0\}$ and $m := -\sum_{i \in N} z_i > 0$. Since $\sum_i z_i = \Delta^+(p) > 0$, we have $\sum_{i:\, z_i > 0} z_i = \Delta^+(p) + m$. Construct $\tilde{z}$ by setting $\tilde{z}_i = 0$ for $i \in N$ and reducing the positive coordinates (each of which has lower bound 0) by total mass m so that $\sum_i \tilde{z}_i = \Delta^+(p)$. This preserves feasibility because it only increases negative coordinates up to 0 and decreases positive coordinates down toward 0. Moreover, $\sum_i \tilde{z}_i^2/(2d_i) < \sum_i z_i^2/(2d_i)$, contradicting optimality. Thus every minimizer must satisfy $z_i \ge 0$ for all i. For indices with $p < p_{0,i}$ we have $\overline{z}_i^+(p) = u_i(p) - q_i^+(p) = 0$, so $z_i = 0$ there. For indices with $p \geq p_{0,i}$ we have $\underline{z}_i^+(p) = \ell_i(p) - q_i^+(p) = 0$ and $\overline{z}_i^+(p) = u_i(p) - \ell_i(p)$. Therefore (*A*3) reduces to

$$\min_{\{z_i\}_{i:\, p \geq p_{0,i}}} \left\{ \sum_{i:\, p \geq p_{0,i}} \frac{z_i^2}{2d_i} \right\} \quad \text{s.t.} \quad \sum_{i:\, p \geq p_{0,i}} z_i = \Delta^+(p), \quad 0 \leq z_i \leq u_i(p) - \ell_i(p),$$

which is exactly $\Psi^+(p)$ when $\Delta^+(p) > 0$ (since then $A^+(p) = \{i\colon p \geq p_{0,i}\}$ and $|\Delta^+(p)| = \Delta^+(p)$). The case $\Delta^+(p) < 0$ is symmetric: any minimizer must satisfy $z_i \leq 0$ for all i; then $z_i = 0$ for all i with $p \geq p_{0,i}$ (because $\underline{z}_i^+(p) = 0$ there), and setting $\delta_i := -z_i \geq 0$ on $\{i\colon p < p_{0,i}\}$ yields exactly the definition of $\Psi^+(p)$ when $\Delta^+(p) < 0$ (since then $A^+(p) = \{i\colon p < p_{0,i}\}$ and $|\Delta^+(p)| = -\Delta^+(p)$). ∎

Since $Z^+(p)$ is a nonempty compact polytope defined by continuous bounds and a continuous affine equality, the correspondence $p \mapsto Z^+(p)$ is continuous. Consequently, by Berge's Maximum theorem $\widetilde{\Psi}^+(\cdot)$ (and so $\Psi^+(\cdot)$) is continuous on $\mathcal{I}$. ∎

Lemma 4. *Posit Assumptions 2 and 3 and fix $p \in \mathcal{I}$. Then, the inner problem of maximizing $\Phi(P)$ over $P \in \mathcal{P}$ satisfying $p^*(P) = p$ is equivalent to minimizing $\int_{p_{0,i}}^{p} q_i$ for $i \in \mathcal{I}_+(p)$ and maximizing $\int_{p}^{p_{0,i}} q_i$ for $i \in \mathcal{I}_-(p)$, subject to feasibility and $\sum_i q_i(p) = \bar{Q}$.*

Proof of Lemma 4. Fix $p \in \mathcal{I}$ and i. Under Assumption 3, the equal-shadow allocation satisfies $0 < q_i^*(P) < q_i^{\max}$, so $P_i\left(q_i^*(P)\right) = p^*$. By the fundamental theorem of calculus,

$$\int_{q_i^{\text{obs}}}^{q_i^*(P)} \left(P_i(x) - p^*\right) dx = q_i^{\text{obs}} \left(p^* - p_{0,i}\right) - \int_{p_{0,i}}^{p^*} q_i(s) ds; \tag{A4}$$

and, summing (A4) over i and using $\sum_i q_i^{\text{obs}} = \bar{Q}$ we obtain

$$\Phi(P) = \bar{Q}p^* - \sum_{i=1}^{n} q_i^{\text{obs}} p_{0,i} - \sum_{i=1}^{n} \int_{p_{0,i}}^{p^*} q_i(s) ds. \tag{A5}$$

Next, as $q_i(p^*) = q_i^*(P)$ and $q_i(p^*) \in [\ell_i(p^*), u_i(p^*)]$ for all i

$$L(p^*) \leq \sum_{i=1}^{n} q_i(p^*) = \sum_{i=1}^{n} q_i^*(P) = \bar{Q} \leq U(p^*) \quad \Longrightarrow \quad p^* \in \mathcal{I}.$$

Observe that

$$-\int_{p_{0,i}}^{p} q_i(s)ds = \begin{cases} -\int_{p_{0,i}}^{p} q_i(s)ds, & i \in \mathcal{I}_+(p), \\ \int_{p}^{p_{0,i}} q_i(s)ds, & i \in \mathcal{I}_-(p). \end{cases}$$

Consequently, conditional on p, maximizing $\Phi(P)$ is equivalent to minimizing $\int_{p_{0,i}}^{p} q_i$ for $i \in \mathcal{I}_+(p)$ and maximizing $\int_{p}^{p_{0,i}} q_i$ for $i \in \mathcal{I}_-(p)$, subject to $\sum_i q_i(p) = \bar{Q}$. ■

Lemma 5. *$q_i'(s) \in [\alpha_i, \beta_i]$ for a.e. $s \in [\min\{p, p_{0,i}\}, \max\{p, p_{0,i}\}]$.*

Proof of Lemma 5. For any $x_1 < x_2$,

$$g_{i,L} \leq \frac{P_i(x_2) - P_i(x_1)}{x_2 - x_1} \leq g_{i,U},$$

and inverting this secant inequality yields

$$\alpha_i \leq \frac{q_i(p_2) - q_i(p_1)}{p_2 - p_1} \leq \beta_i \quad \text{for all } p_1 < p_2 \text{ in the relevant range.}$$

Take limits along the differentiability points of the Lipschitz function q_i. ■

Lemma 6. *Now let $t_0 < t_1$, $\alpha < \beta$, and let function q be absolutely continuous on $[t_0, t_1]$ with $q(t_0) = q_0$ and $q'(t) \in [\alpha, \beta]$ a.e. Define $\underline{q}(t) := q_0 + \alpha(t - t_0)$. If $q(t_1) = \underline{q}(t_1) + \delta$ for some $\delta \in [0, (\beta - \alpha)(t_1 - t_0)]$, then*

$$\int_{t_0}^{t_1} q(t)dt \geq \int_{t_0}^{t_1} \underline{q}(t)dt + \frac{\delta^2}{2(\beta - \alpha)}. \tag{A6}$$

Moreover, equality holds if and only if q is "bang-bang;" $q'(t) = \alpha$ a.e. on $[t_0, t_1 - h]$ and $q'(t) = \beta$ a.e. on $[t_1 - h, t_1]$, where $h = \delta / (\beta - \alpha)$.

Proof of Lemma 6. Write $q'(t) = \alpha + u(t)$ with $u(t) \in [0, \beta - \alpha]$ a.e. on $[t_0, t_1]$. The two constraints on u are:

$$0 \leq u(t) \leq \beta - \alpha \quad \text{for a.e. } t \in [t_0, t_1] \qquad \text{(box constraint)},$$

and

$$\int_{t_0}^{t_1} u(t)dt = \delta \qquad \text{(total mass constraint)}.$$

Let $\underline{q}(t) := q_0 + \alpha(t - t_0)$. Then

$$q(t) = \underline{q}(t) + \int_{t_0}^{t} u(s)ds.$$

Integrating over $[t_0, t_1]$ and using Fubini gives

$$\int_{t_0}^{t_1} q(t)dt = \int_{t_0}^{t_1} \underline{q}(t)dt + \int_{t_0}^{t_1} \left(\int_{t_0}^{t} u(s)ds \right) dt = \int_{t_0}^{t_1} \underline{q}(t)dt + \int_{t_0}^{t_1} (t_1 - s)u(s)ds.$$

Thus, conditional on δ, minimizing $\int_{t_0}^{t_1} q(t)dt$ is equivalent to minimizing

$$\min_{u} \int_{t_0}^{t_1} (t_1 - s)u(s)ds \quad \text{s.t.} \quad 0 \le u(s) \le \beta - \alpha \text{ a.e.}, \int_{t_0}^{t_1} u(s)ds = \delta. \tag{A7}$$

Define the cumulative mass $U(t) := \int_{t_0}^{t} u(s)ds$. U is absolutely continuous, nondecreasing, satisfies $U(t_0) = 0$, $U(t_1) = \delta$, and $U'(t) = u(t) \in [0, \beta - \alpha]$ a.e. Integrating by parts:

$$\int_{t_0}^{t_1} (t_1 - s)u(s)ds = \int_{t_0}^{t_1} (t_1 - s)U'(s)ds = [(t_1 - s)U(s)]_{t_0}^{t_1} - \int_{t_0}^{t_1} (-1)U(s)ds = \int_{t_0}^{t_1} U(s)ds.$$

So the program in ($A7$) is equivalent to

$$\min_{U} \int_{t_0}^{t_1} U(t)dt \quad \text{s.t.} \quad U(t_0) = 0,\ U(t_1) = \delta,\ 0 \le U'(t) \le \beta - \alpha \text{ a.e.}$$

Let $h := \delta/(\beta - \alpha) \in [0, t_1 - t_0]$. Since $U'(s) \le \beta - \alpha$ a.e., for any $t \in [t_0, t_1]$,

$$\delta - U(t) = U(t_1) - U(t) = \int_{t}^{t_1} U'(s)ds \le \int_{t}^{t_1} (\beta - \alpha)ds = (\beta - \alpha)(t_1 - t),$$

hence,

$$U(t) \ge \delta - (\beta - \alpha)(t_1 - t) = (\beta - \alpha)\left(t - (t_1 - h)\right).$$

Also $U(t) \ge 0$. Therefore, for all $t \in [t_0, t_1]$,

$$U(t) \ge U^*(t) := \max\left\{0, (\beta - \alpha)\left(t - (t_1 - h)\right)\right\}.$$

Integrating both sides yields

$$\int_{t_0}^{t_1} U(t)dt \geq \int_{t_0}^{t_1} U^*(t)dt = \int_{t_1-h}^{t_1} (\beta - \alpha)\,(t - (t_1 - h))\,dt = \frac{(\beta-\alpha)h^2}{2} = \frac{\delta^2}{2(\beta-\alpha)}.$$

Moreover, equality holds if and only if $U(t) = U^*(t)$ for all t, which forces $U'(t) = 0$ a.e. on $[t_0, t_1 - h]$ and $U'(t) = \beta - \alpha$ a.e. on $[t_1 - h, t_1]$, i.e.

$$u(t) = \begin{cases} 0, & t \in [t_0, t_1 - h] \text{ a.e.}, \\ \beta - \alpha, & t \in [t_1 - h, t_1] \text{ a.e.} \end{cases}$$

That is, the optimizer concentrates all mass δ on the terminal interval $[t_1 - h, t_1]$, saturating the box constraint. This is a special case of the bathtub principle (the Hardy-Littlewood rearrangement) for minimizing a linear functional under an integral constraint. ■

Claim 6 *Fix $p \in \mathcal{I}$.*

I. *If $\Delta^+(p) > 0$, there exist numbers $\{\delta_i\}_{i\in\mathcal{I}_+(p)}$ such that $0 \leq \delta_i \leq u_i(p) - \ell_i(p)$ for all $i \in \mathcal{I}_+(p)$ and $\sum_{i\in\mathcal{I}_+(p)} \delta_i = \Delta^+(p)$, and the endpoint vector $q(p)$ defined by*

$$q_i(p) = \ell_i(p) + \delta_i \text{ for } i \in \mathcal{I}_+(p), \qquad q_i(p) = u_i(p) \text{ for } i \in \mathcal{I}_-(p)$$

satisfies $\sum_{i=1}^n q_i(p) = \bar{Q}$. Moreover, in the inner problem, there is an optimizer with $q_i(p) = u_i(p)$ for all $i \in \mathcal{I}_-(p)$; equivalently, deviations from the baseline endpoints $q_i^+(p)$ may be taken without loss of optimality to occur only on $\mathcal{A}^+(p) = \mathcal{I}_+(p)$.

II. *If $\Delta^+(p) < 0$, there exist numbers $\{\delta_i\}_{i\in\mathcal{I}_-(p)}$ such that $0 \leq \delta_i \leq u_i(p) - \ell_i(p)$ for all $i \in \mathcal{I}_-(p)$ and $\sum_{i\in\mathcal{I}_-(p)} \delta_i = -\Delta^+(p)$, and the endpoint vector $q(p)$ defined by*

$$q_i(p) = \ell_i(p) \text{ for } i \in \mathcal{I}_+(p), \qquad q_i(p) = u_i(p) - \delta_i \text{ for } i \in \mathcal{I}_-(p)$$

satisfies $\sum_{i=1}^n q_i(p) = \bar{Q}$. Moreover, in the inner problem, there is an optimizer with $q_i(p) = \ell_i(p)$ for all $i \in \mathcal{I}_+(p)$; equivalently, deviations from the baseline endpoints $q_i^+(p)$ may be taken without loss of optimality to occur only on $\mathcal{A}^+(p) = \mathcal{I}_-(p)$.

Proof of Claim 6. We prove (i) as the proof of (ii) is identical with signs reversed.

Assume $\Delta^+(p) > 0$. Since $p \in \mathcal{I}$, we have $\bar{Q} \le U(p) = \sum_{i=1}^n u_i(p)$. Using

$$\sum_{i=1}^n q_i^+(p) = \sum_{i \in \mathcal{I}_+(p)} \ell_i(p) + \sum_{i \in \mathcal{I}_-(p)} u_i(p),$$

we obtain

$$\Delta^+(p) = \bar{Q} - \sum_{i \in \mathcal{I}_+(p)} \ell_i(p) - \sum_{i \in \mathcal{I}_-(p)} u_i(p) \le \sum_{i \in \mathcal{I}_+(p)} \left(u_i(p) - \ell_i(p)\right).$$

Thus the total upward slack available in the coordinates $i \in \mathcal{I}_+(p)$ is at least $\Delta^+(p)$, so we can choose $\{\delta_i\}_{i \in \mathcal{I}_+(p)}$ with $0 \le \delta_i \le u_i(p) - \ell_i(p)$ and $\sum_{i \in \mathcal{I}_+(p)} \delta_i = \Delta^+(p)$. Defining $q(p)$ as in the statement then yields $\sum_i q_i(p) = \bar{Q}$ by construction.

For the "without loss" part, let $\{q_i(\cdot)\}_{i=1}^n$ be an optimizer for the inner problem at this p. Conditional on p, the objective depends on $q_i(\cdot)$ only through the signed integrals

$$-\int_{p_{0,i}}^{p} q_i(s)ds \text{ for } i \in \mathcal{I}_+(p), \qquad \text{and} \quad \int_{p}^{p_{0,i}} q_i(s)ds \text{ for } i \in \mathcal{I}_-(p).$$

By Lemmas 5 and 6, for each market k and each admissible endpoint value $q_k(p) \in [\ell_k(p), u_k(p)]$, the choice of $q_k(\cdot)$ that optimizes its signed integral term is the Lemma 6 equality-case ("bang-bang") form. Moreover, the resulting value depends on the endpoint only through a quadratic penalty in the endpoint gap: it is strictly decreasing in $q_k(p) - \ell_k(p)$ when $k \in \mathcal{I}_+(p)$, and strictly decreasing in $u_k(p) - q_k(p)$ when $k \in \mathcal{I}_-(p)$.

If an optimizer had some $i \in \mathcal{I}_-(p)$ with $q_i(p) < u_i(p)$, then (because $\sum_{k=1}^n q_k(p) = \bar{Q} > \sum_{k=1}^n q_k^+(p)$) there must exist some $j \in \mathcal{I}_+(p)$ with $q_j(p) > \ell_j(p)$. Take

$$\varepsilon \in \left(0, \min\left\{u_i(p) - q_i(p),\, q_j(p) - \ell_j(p)\right\}\right],$$

and modify only these two endpoints by setting

$$\tilde{q}_i(p) := q_i(p) + \varepsilon, \qquad \tilde{q}_j(p) := q_j(p) - \varepsilon,$$

leaving $\tilde{q}_k(p) := q_k(p)$ for all $k \notin \{i, j\}$. This preserves $\sum_{k=1}^n \tilde{q}_k(p) = \bar{Q}$ and the box constraints $\tilde{q}_k(p) \in [\ell_k(p), u_k(p)]$. Now choose for markets i and j the Lemma 6 equality-case forms under the modified endpoints (leaving all other markets unchanged). Since the ε-transfer strictly reduces both endpoint gaps $u_i(p) - q_i(p)$ and $q_j(p) - \ell_j(p)$, it strictly reduces the associated quadratic penalties and therefore strictly increases the objective, contradicting optimality. Hence, every optimizer must satisfy $q_i(p) = u_i(p)$ for all $i \in \mathcal{I}_-(p)$, so deviations from $q_i^+(p)$ may be taken to occur only on $\mathcal{I}_+(p) = \mathcal{A}^+(p)$. ■

References

AAA. 1974. *AAA Fuel Gauge Report II.* Technical report 5. Falls Church, VA, February.

Akbarpour, Mohammad, Piotr Dworczak, and Scott Duke Kominers. 2024. "Redistributive Allocation Mechanisms." *Journal of Political Economy* 132 (6): 1831–1875.

Associated Press. 1973. "Chicks Smothered as Farmers Decry Price Freeze." *The New York Times* (June). ISSN: 0362-4331, accessed March 19, 2025. https://www.nytimes.com/1973/06/27/archives/chicks-smothered-as-farmers-decry-price-freeze.html.

Barzel, Yoram. 1974. "A Theory of Rationing by Waiting." *Journal of Law and Economics* 17 (1): 73–95.

Bradley, Robert L. 1997. *Oil, Gas and Government: The U.S. Experience.* Rowman & Littlefield Inc, October. ISBN: 0-8476-8108-4.

Brunt, Christopher S. 2025. "Moving beyond Harberger's Triangle to present the inefficiency from misallocated market transactions under price controls" [in en]. *The Journal of Economic Education* 56, no. 4 (October): 322–328. ISSN: 0022-0485, 2152-4068, accessed December 26, 2025. https://doi.org/10.1080/00220485.2025.2524631. https://www.tandfonline.com/doi/full/10.1080/00220485.2025.2524631.

Bulow, Jeremy I., and Paul D. Klemperer. 2012. "Regulated Prices, Rent Seeking, and Consumer Surplus." *Journal of Political Economy* 120 (1): 160–186. Accessed March 9, 2018.

Buurma-Olsen, Jennifer, Hans R. A. Koster, Jos van Ommeren, and Jort Sinninghe Damsté. 2025. "Quantifying misallocation of public housing." *Journal of Public Economics* 242 (February): 105272. ISSN: 0047-2727, accessed January 14, 2026. https://doi.org/10.1016/j.jpubeco.2024.105272. https://www.sciencedirect.com/science/article/pii/S0047272724002081.

Condorelli, Daniele. 2013. "Market and Non-Market Mechanisms for the Optimal Allocation of Scarce Resources." *Games and Economic Behavior* 82:582–591.

Davis, Lucas W., and Lutz Kilian. 2011. "The Allocative Cost of Price Ceilings in the U.S. Residential Market for Natural Gas" [in en]. *Journal of Political Economy* 119, no. 2 (April): 212–241. ISSN: 0022-3808, 1537-534X, accessed May 15, 2025. https://doi.org/10.1086/660124. https://www.journals.uchicago.edu/doi/10.1086/660124.

Deacon, Robert T., Walter J. Mead, and Vinod B. Agarwal. 1980. "Price Controls and International Petroleum Product Prices." Publisher: Department of Energy.

Deacon, Robert T., and Jon Sonstelie. 1985. "Rationing by Waiting and the Value of Time: Results from a Natural Experiment." *Journal of Political Economy* 93 (4): 627–647.

———. 1989. "The Welfare Costs of Rationing by Waiting." *Economic Inquiry* 27:179–196.

Diamond, Rebecca, Tim McQuade, and Franklin Qian. 2019. "The Effects of Rent Control Expansion on Tenants, Landlords, and Inequality: Evidence from San Francisco." *American Economic Review* 109 (9): 3365–3394. https://doi.org/10.1257/aer.20181289.

Federal Highway Administration. 2012. *Highway Statistics, Table MF-202: Motor-Fuel Use.* Annual net volume of gasoline taxed; used as a gasoline-specific cross-check on 1974 trend shortfall. https://www.fhwa.dot.gov/policyinformation/statistics/2012/mf202.cfm.

Filtzer, Donald. 1996. "Labor Discipline, the Use of Work Time, and the Decline of the Soviet System, 1928-1991." *International Labor and Working-Class History*, no. 50, 9–28. ISSN: 0147-5479, accessed February 3, 2026. https://www.jstor.org/stable/27672305.

Frech, H. E., III, and William C. Lee. 1987. "The Welfare Cost of Rationing-By-Queuing Across Markets: Theory and Estimates from the U. S. Gasoline Crises*." *The Quarterly Journal of Economics* 102, no. 1 (February): 97–108. ISSN: 0033-5533, accessed April 23, 2025. https://doi.org/10.2307/1884682. https://doi.org/10.2307/1884682.

Glaeser, Edward L., and Erzo F. P. Luttmer. 2003. "The Misallocation of Housing under Rent Control." *American Economic Review* 93 (4): 1027–1046. https://doi.org/10.1257/000282803769206188.

Gordon, Robert J. 1973. "The Response of Wages and Prices to the First Two Years of Controls." *Brookings Papers on Economic Activity* 1973 (3): 765–815.

Grünbaum, Branko. 2003. *Convex Polytopes.* 2nd. Vol. 221. Graduate Texts in Mathematics. New York: Springer.

Gupte, Pranay B. 1979. "FUEL CRISIS TERMED CRITICAL AT PUMPS IN NEW YORK REGION." *The New York Times* (June). ISSN: 0362-4331, accessed February 7, 2026. https://www.nytimes.com/1979/06/25/archives/fuel-crisis-termed-critical-at-pumps-in-new-york-region-long-lines.html.

Harberger, Arnold C. 1954. "Monopoly and Resource Allocation." *American Economic Review* 44 (2): 77–87.

Hausman, Jerry A., and Whitney K. Newey. 1995. "Nonparametric Estimation of Exact Consumers Surplus and Deadweight Loss." *Econometrica* 63 (6): 1445–1476.

———. 2016. "Individual Heterogeneity and Average Welfare." *Econometrica* 84 (3): 1225–1248.

Hughes, Jonathan E., Christopher R. Knittel, and Daniel Sperling. 2008. "Evidence of a Shift in the Short-Run Price Elasticity of Gasoline Demand." *The Energy Journal* 29 (1): 113–134.

Kang, Zi Yang. 2026. "Robust Regulation: Prices vs. Quantities." Working paper, University of Toronto.

Kang, Zi Yang, and Shoshana Vasserman. 2025. "Robustness Measures for Welfare Analysis" [in en]. *American Economic Review* 115, no. 8 (August): 2449–2487. ISSN: 0002-8282, accessed January 14, 2026. https://doi.org/10.1257/aer.20220673. https://www.aeaweb.org/articles?id=10.1257/aer.20220673.

Kholodilin, Konstantin A. 2024. "Rent control effects through the lens of empirical research: An almost complete review of the literature." *Journal of Housing Economics* 63:101976.

Kominers, Scott Duke, and Piotr Dworczak. 2025. "A Price Theory of Price Gouging."

Kornai, János. 1992. *The Socialist System: The Political Economy of Communism.* Princeton, NJ: Princeton University Press. ISBN: 978-0-691-00393-1.

Kremer, Michael, and Christopher M. Snyder. 2018. *Worst-Case Bounds on R&D and Pricing Distortions.* Working Paper 25119. NBER.

Mulligan, Casey B. 2025. "Equilibrium Responses to Price Controls: A Supply-Chain Approach." *Public Choice* 203:23–52. https://doi.org/10.1007/s11127-024-01196-8.

Murphy, Kevin M., Andrei Shleifer, and Robert W. Vishny. 1992. "The Transition to a Market Economy: Pitfalls of Partial Reform." Publisher: Oxford University Press, *The Quarterly Journal of Economics* 107, no. 3 (August): 889–906. Accessed October 31, 2017. https://doi.org/10.2307/2118367.

Office, Federal Energy. 1974. *Review of Complaints Concerning the Mandatory Petroleum Allocation Program And The Regulation of Petroleum Pricing.* Technical report 090294. https://www.gao.gov/assets/b-178205-090294.pdf.

Oi, Walter Y. 1976. "On Measuring the Impact of Wage-Price Controls: A Critical Appraisal." *Carnegie-Rochester Conference Series on Public Policy* 2:7–64. https://doi.org/10.1016/S0167-2231(76)80003-6.

Poole, Robert W. 1973. "Fuel Shortage: A Prehistory." *Reason.*

Rockoff, Hugh. 1984. *Drastic Measures: A History of Wage and Price Controls in the United States.* New York: Cambridge University Press.

Smith, Hedrick. 1984. *The Russians.* Google-Books-ID: LFpbPC2k7NgC. Ballantine Books. ISBN: 978-0-345-31746-9.

Sowell, Thomas. 1980. *Knowledge and Decisions.* New York: BasicBooks. Reissued 1996.

The New York Times. 1973. "Baby Chicks Killed and Cooked for Feed." *The New York Times* (June). ISSN: 0362-4331, accessed February 1, 2026. https://www.nytimes.com/1973/06/25/archives/baby-chicks-killed-and-cooked-for-feed.html.

TIME. 1973. *CONTROLS: A Threat of Food Shortage,* July. Accessed March 19, 2025. https://time.com/archive/6841371/controls-a-threat-of-food-shortage/.

TIME. 1974. *SHORTAGES: Gas Fever: Happiness Is a Full Tank,* February. Accessed February 7, 2026. https://time.com/archive/6875854/shortages-gas-fever-happiness-is-a-full-tank/.

Turkel, Tux. 2023. *Fifty years ago, the energy crisis officially began. And it hasn't ended.,* September. Accessed February 10, 2026. https://themainemonitor.org/fifty-years-ago-the-climate-crisis-began-and-it-hasnt-ended/.

U.S. Congress, Joint Economic Committee, Subcommittee on Consumer Economics. 1973. *The gasoline and fuel oil shortage: Hearings before the subcommittee on consumer economics of the joint economic committee, congress of the united states, ninety-third congress, first session, may 1, 2, and june 2, 1973.* Number: 99-740 O tex.entrytype: government. Washington, DC.

U.S. Energy Information Administration. 2025. *Monthly Energy Review, Table 3.1: Petroleum Overview, Annual (1949–2024).* Accessed via Monthly Energy Review tables; products supplied series used for pre-shock trend comparison. https://www.eia.gov/totalenergy/data/monthly/index.php.

Verleger, Philip K. 1979. "The U.S. Petroleum Crisis of 1979." *Brookings Papers on Economic Activity* 1979 (2): 463–476.

Weitzman, Martin L. 1977. "Is the Price System or Rationing More Effective in Getting a Commodity to Those Who Need It Most?" *Bell Journal of Economics* 8:517–524.

Yergin, Daniel. 1991. *The prize: the epic quest for oil, money, and power* [in eng]. New York: Simon & Schuster. ISBN: 978-0-671-50248-5.