

Pricing in Crisis

Reyer Gerlagh Matti Liski Iivo Vehviläinen

May 21, 2026

Abstract

When demand aggregates both price-sensitive and price-insensitive behaviors, uniform pricing becomes a deficient market design that generates negative surplus during extreme-price events. We develop a price-control mechanism that efficiently resolves the tradeoff between protecting consumers and limiting rents. The mechanism implements a dynamic price cap that responds to demand adjustments and induces truthful supply through incentive payments. In a quantitative application to the French wholesale electricity market during the 2022–2023 energy crisis, the mechanism would have lowered expected procurement costs by roughly €200 billion, about two-thirds of total projected costs in this central scenario.

1 Introduction

Markets in which a homogeneous good is traded at a uniform price often aggregate buyers with heterogeneous abilities to adjust their demand in response to price changes. Some buyers can flexibly reduce or shift their quantities when prices fluctuate, while others cannot, either because of technological rigidities, contractual obligations, or the essential nature of the good. This mixture of elastic and inelastic demands is a

Gerlagh (r.gerlagh@uvt.nl): Economics department of the Tilburg University. Liski (matti.liski@aalto.fi) and Iivo Vehviläinen (iivo.vehvilainen@aalto.fi): Economics department of the Aalto University, and Helsinki GSE. This project has been funded by Research Council of Finland COP27 (project 357065) and the Finnish Cultural Foundation.

structural feature of many high-stakes markets. When supply is tight, the clearing price can rise to extreme levels because aggregate demand reflects the willingness to pay only of those who are able to adjust.

From a welfare perspective, this creates a fundamental tension. Leaving prices unconstrained exposes inelastic buyers to shocks, leading to large transfers and inefficient destruction of surplus, while constraining prices too much eliminates incentives for efficient allocation on the responsive margin of supply and demand. Because discrimination between demand types is often infeasible or undesirable, policies such as price control and rationing cannot be targeted to specific groups. The central design challenge is therefore to balance protection of inelastic demand with the preservation of efficiency on the responsive margin.

We start by noting that this problem has, in principle, a simple structure. For any realized demand state (i.e., demand schedule), there exists a unique price threshold at which total consumer surplus is zero, aggregating both the inelastic and elastic components. Under standard regularity conditions, capping the price at this threshold and rationing total demand to clear the market maximizes total surplus. Crucially, this threshold depends only on the composition of demand, not on the shape or level of supply. It can therefore be computed directly whenever the demand schedule is observable and decomposable into responsive and inelastic components.

This result implies a transparent rule for price control: the optimal price cap adjusts dynamically with changes in demand composition, tightening when inelastic demand dominates and relaxing when elastic demand expands. Such a rule cushions price shocks without suppressing the information contained in prices and preserves incentives for efficient adjustment.

Our analysis is cast in general terms, taking as primitives demand and supply schedules as in market exchanges. Electricity markets provide a leading example: our main quantitative application is to the European energy crisis of 2022-2023, discussed below. But price-control elements are in everyday use also in other high-stakes markets, including short-term money markets. In these markets, a significant share of liquidity demand is effectively committed in the short run—arising from payment obligations, collateral or margin calls, and reserve requirements—while the remainder adjusts elastically to interest rates. Short-term money markets are a central venue in which policy rates are implemented.¹

¹In corridor systems, policy implementation relies on administered lending and deposit rates that

This paper makes the following contributions. First, we characterize the welfare-maximizing price cap in uniform-price markets with elastic and inelastic demand and provide an algorithm to compute it from observed demand schedules. Second, we provide a mechanism-design foundation for the price-cap algorithm by developing an auction mechanism that induces truthful supply behavior. Third, we allow for heterogeneity in buyers’ welfare weights and endogenous selection into price responsiveness and show that the price-cap mechanism remains implementable and optimal. Fourth, we apply the theory to the European wholesale electricity market during the 2022 energy crisis, focusing on France, and quantify its implications.

Empirical case. In fall 2022, expectations for 2023 electricity procurement costs escalated—by roughly a factor of ten in France—as reflected in forward contract prices. Using hourly market outcomes for 2023 in France, we quantify which factors can rationalize this surge. We construct counterfactual supply curves that isolate three channels: input costs (notably gas), nuclear availability, and the EU price-capping protocol in place at the time. The results show that reduced nuclear availability, interacting with the prevailing EU protocol, is sufficient to account for the observed escalation, whereas input costs alone cannot. Replacing the EU protocol with our price-control mechanism would have reduced expected procurement costs by about two thirds (roughly €200 billion), substantially mitigating the threat to market stability. While rationing is more frequent under the optimal policy, the total quantity rationed is not materially higher than what market already expected under the EU protocol. We also compute the incentive payments required to sustain truthful supply of marginal capacity. Under the optimal price-cap policy, these payments fall sharply—to only a few euros per MWh—and are fiscally negligible relative to suppliers’ revenues from sales at market prices.

Related literature. Our work contributes to several strands of the literature.

A first strand concerns price controls and rationing as exogenous institutions. A classic contribution is Weitzman (1977), showing how rationing can improve allocation

bound short-term money-market rates; see e.g. Berentsen and Monnet (2008) and Whitesell (2006). In practice, these price instruments are embedded in operational frameworks that can also specify allotment procedures, counterparty access, collateral eligibility, and haircuts (e.g., (Ashcraft et al., 2011; De Fiore et al., 2018)). Välimäki (2003) reviews and analyzes fixed-rate tender experiences in which administered rates are combined with proportional rationing.

when some buyers cannot afford essential goods (see also Sah, 1987; Wijkander, 1988). Other studies highlight the distributional and efficiency properties of rationing in the management of common resources (e.g., Donna and Espin-Sanchez, 2023; Ryan and Sudarshan, 2022). In contrast to this literature, which typically treats rationing or price control as exogenous, our result shows that a rationing rule—implemented together with an optimal price cap—arises endogenously as a welfare-maximizing policy even in the absence of distributional weights or externalities. This result reflects the presence of a committed component of demand, which may itself be technologically or institutionally determined, and contrasts with the large price-theory literature emphasizing the efficiency losses from rationing (e.g., Bulow and Klemperer, 2012). The key novelty is that the optimal rationing rule and price cap emerge directly from demand primitives, without invoking additional policy objectives or institutional assumptions.

A second strand of the literature concerns price formation and scarcity pricing in electricity markets. These analyses emphasize that when some consumers cannot respond to price changes, efficiency requires treating demand types differently: inelastic consumers should face fixed prices and be rationed when needed, while elastic consumers should continue to face market prices that guide efficient adjustment (e.g., Wilson, 1989; Joskow and Tirole, 2006, 2007; Borenstein and Holland, 2005; Gowrisankaran et al., 2016). In practice, however, electricity *wholesale* markets clear all demand, derived from different sources, at a single uniform price and do not distinguish between demand types at the market-clearing stage. Our planner would like to follow the explicit discrimination policy but must work with aggregate demand to infer the demand composition for obtaining the welfare-maximizing price cap.²

A third literature concerns mechanisms that address market power and informational distortions on the supply side. Within a price-theory framework, efficient allocations can be implemented by clearing the market $n + 1$ times based on the submitted bids. The n counterfactual clearings generate discriminatory unit prices that effectively transform the setting into a Vickrey auction (Vickrey, 1961) with VCG payments, as in Montero (2008). This avoids the inefficiencies of pay-as-bid auctions (cf. Ausubel et al., 2014). Multi-unit Vickrey auctions are uncommon in practice (cf.

²The literature on activating consumer price-sensitivity identifies a related but distinct distortion (Fowle et al., 2021; Ito et al., 2023). Our mechanism remains optimal even after correcting for possible behavioral biases in consumer selection into price-sensitivity; we assume no behavioral biases but only differing private costs of technologies.

Hortaçsu and McAdams, 2018). In our setting, however, the approach is appealing when combined with a price cap. When the cap binds, firms receive no incentive payments beyond their market revenue, so the mechanism limits rents precisely in scarcity events, in contrast to designs without a cap. As a result, the VCG transfers required to sustain truthful supply are lower. The resulting design optimally resolves the trade-off between shielding buyers from extreme prices and limiting producers' rents.

A fourth strand studies mechanism design under distributional objectives. Recent contributions have shown how welfare weights shape optimal market designs (e.g., Dworzak et al., 2021; Akbarpour et al., 2024; Pai and Strack, 2022; Tokarski et al., 2023). We allow for heterogeneity in buyers' value of money and for endogenous selection into price responsiveness. In this setting, welfare weights shift the optimal price cap downward when the planner places higher value on inelastic buyers, while responsiveness externalities raise the cap. Despite these additional forces, the mechanism remains tractable and implements the constrained-efficient allocation.

2 Optimal Price Control in a Competitive Market

2.1 Environment

Demand and Supply. The policymaker observes the aggregate demand and supply schedules prior to market clearing. It is useful to think of these schedules as bids submitted to an exchange where the policy is implemented. Let $D_x(p)$ denote aggregate demand in demand state x , differentiable and strictly decreasing in price p , and let $S_y(p)$ denote aggregate supply in supply state y , differentiable and strictly increasing in p .³

The joint market state $z = (x, y)$ is drawn from a commonly known joint distribution. The demand function $D_x(p)$ satisfies decreasing differences in (p, x) : a higher price reduces demand more in a higher-demand state. The supply function $S_y(p)$ satisfies increasing differences in (p, y) : a higher price increases supply more in a higher-supply state. The model therefore allows for shocks to demand, supply, or both.

³We use differential methods for the analytical results; the empirical application implements these objects as step-functions, as observed in the data.

The policymaker does not need to know the functional dependence of these schedules on the underlying states x and y ; it is sufficient that the realized schedules are observable as data containing quantities demanded and supplied at each price. We may think of each market opening as a new draw of z and its corresponding schedules. The policy described below specifies how outcomes are chosen in each such realization.

The policymaker controls a single instrument, a *rationing rule* $\mu_z \in [0,1]$, that scales the aggregate demand schedule:

$$D_z(p) = \mu_z D_x(p). \tag{2.1}$$

Given $D_z(p)$ and $S_y(p)$, the market clears at the equilibrium price p_z satisfying

$$D_z(p_z) = S_y(p_z).$$

We assume that this equation defines a unique equilibrium price $p_z > 0$ for every state z , and that the expected values $\bar{p} = \mathbb{E}[p_z]$ and $\bar{p}_x = \mathbb{E}[p_z | x]$ remain bounded for all admissible choices by the agents.⁴

Sticky and Non-sticky Demands. Consumers differ in how their individual demand responds to market prices. A share $\theta \in [0,1]$ of consumers are *responsive*, and the remaining $1 - \theta$ are *sticky*.⁵ When responsive, a consumer's private demand is $d_x(p)$, a strictly decreasing and differentiable function of price p in any demand state x . When sticky, the consumer does not observe the realized market price but forms rational expectations \bar{p}_x . Sticky demand is fixed at the expected-price level $d_x(\bar{p}_x)$, denoted \bar{d}_x for short. Sticky consumers thus understand how their preset demand varies across states and how it correlates with prices, even though it does not adjust to the realized price.⁶

The aggregate demand in any state x can then be written as

$$D_x(p) = (1 - \theta) \bar{d}_x + \theta d_x(p), \tag{2.2}$$

⁴We omit the subscript z from expectations when no confusion arises.

⁵Section 4 endogenizes the self-selection into responsiveness.

⁶We collapse the indexation of realized demand schedules D_x and price expectations \bar{p}_x . More generally, one can let x determine D_x while forming expectations with respect to an information set X_x such that $x \in X_x$, so that $\bar{p}_x = \mathbb{E}[p_z | X_x]$. We suppress this distinction.

where θ is the share of responsive consumers. The policymaker can identify the two components, \bar{d}_x and $d_x(p)$, from observed aggregate data $D_x(p)$ if information on \bar{p}_x is available—for instance, from contract prices or from expert evaluations linking expected prices to demand covariates.

Given $D_x(p)$ and \bar{p}_x , identification of sticky demand follows from $D_x(\bar{p}_x) = d_x(\bar{p}_x)$, which pins down \bar{d}_x . The share $\theta > 0$ can be inferred from $\lim_{p \rightarrow \infty} D_x(p) = (1 - \theta) \bar{d}_x$, and the responsive demand from

$$d_x(p) = \frac{D_x(p) - (1 - \theta) \bar{d}_x}{\theta}.$$

With this information, the policymaker can evaluate welfare effects of alternative policies.

To ensure that the welfare evaluations are well defined, we assume that the elasticity of the responsive demand satisfies $\sigma(d_x) > 1$ for sufficiently high p in every state x . This guarantees that the indirect utility from responsive consumption, $u_x(d)$, is finite at $d = 0$. Normalizing $u_x(0) = 0$, utility can be written as

$$u_x(d_x(p)) = p d_x(p) + \int_{p' \geq p} d_x(p') dp'.$$

Understanding how the concept of utility follows from observables, we move on to the planner's problem.

2.2 Planner's Problem

The planner observes the aggregate demand and supply schedules before market clearing and can influence the market outcome through a rationing rule $\mu_z \in [0, 1]$. We also introduce a per-unit tax or subsidy τ_z on consumption, an auxiliary instrument that aids interpretation. Formally, for a given τ_z , we continue to denote the demand by $D_z(p) = \mu_z D_x(p + \tau_z)$.

Aggregate utility in state z is given by

$$U_z = \mu_z [(1 - \theta) u_x(\bar{d}_x) + \theta u_x(d_x)], \tag{2.3}$$

where $d_x = d_x(p_z + \tau_z)$ is the demand of non-sticky consumers. Because rationing applies uniformly, total utility scales proportionally with μ_z .

Aggregate welfare is

$$W_z = U_z - C_y(D_z), \quad (2.4)$$

where $C_y(D_z)$ is the total cost of supplying quantity D_z , strictly convex and differentiable.

2.3 Main Characterization

Using $d_x = [D_z/\mu_z - (1 - \theta)\bar{d}_x]/\theta$, substitute into (2.4) to express welfare as a function of (D_z, μ_z, τ_z) . Differentiating yields

$$dW_z = [u'_x - C'_y] dD_z + \left[\frac{U_z}{\mu_z} - \frac{D_z}{\mu_z} u'_x \right] d\mu_z, \quad (2.5)$$

where u'_x is the marginal utility of non-sticky consumers, evaluated at $d_x(p_z + \tau_z)$, and C'_y is the marginal cost of total supply D_z .

Equation (2.5) shows how welfare changes with aggregate demand and rationing. The planner can vary demand dD_z through τ_z . Trading sets $u'_x - C'_y = \tau_z$ and it follows immediately that it is not optimal to distort the choices at the adjusting margin by a consumption tax or subsidy.

Proposition 2.1 (Optimal Policy). *The optimal policy does not use taxes or subsidies, $\tau_z = 0$, and regulates the market solely through rationing μ_z . For every state z , the optimal pair (p_z^*, μ_z^*) satisfies*

$$p_z D_z(p_z) \leq U_z \perp \mu_z \leq 1. \quad (2.6)$$

That is, rationing applies if and only if consumer surplus is zero at the equilibrium price.

Proof. The first term in (2.5) shows that introducing a tax τ_z would create a gap between u'_x and C'_y , lowering welfare. The first term thus shows that $\tau_z = 0$. The market equilibrium adjusts to variations in μ_z . The second term implies that welfare increases with μ_z when $\frac{U_z}{\mu_z} - \frac{D_z}{\mu_z} u'_x > 0$ and decreases otherwise, establishing the complementarity condition in (2.6). \square

Condition (2.6) defines a state-contingent price boundary p_z^* , implemented as a price cap. Because both U_z and D_z scale proportionally with μ_z , the boundary depends

only on the demand state x , denoted p_x^* . Thus, remarkably, the optimal policy implies that the planner needs only information about demand to derive the price boundary. When binding, the price cap equates the average utility per unit of consumption,

$$p_x^* = \frac{(1 - \theta) u_x(\bar{d}_x) + \theta u_x(d_x^*)}{(1 - \theta) \bar{d}_x + \theta d_x^*}, \quad (2.7)$$

where $d_x^* = d_x(p_x^*)$. Intuitively, the price p_x^* prevents a negative surplus by balancing the positive consumer surplus from non-sticky consumers and negative surplus from sticky ones.

Implementation of the Optimal Policy. We can now state the rationing policy in full:

Proposition 2.2. *Assume that the demand and supply schedules, $D_x(p)$ and $S_y(p)$, are observable. Then:*

- (i) *For each demand state x , the unique optimal price cap p_x^* derived from $D_x(p)$ satisfies (2.7) with $d_x^* = d_x(p_x^*)$, and the relevant utility terms are identified from the observed $D_x(p)$.*
- (ii) *For each x , there exists a unique supply state y_x^* such that*

$$S_{y_x^*}(p_x^*) = (1 - \theta) \bar{d}_x + \theta d_x^*.$$

If and only if $y < y_x^$, it is optimal to implement a binding price cap at $p = p_x^*$ with demand rationed by*

$$\mu_z^* = \frac{S_y(p_x^*)}{D_x(p_x^*)}.$$

The optimal rationing loosens, and the rationing price increases, with the share of non-sticky consumers:

$$\frac{\partial \mu_z^*}{\partial \theta} > 0, \quad \frac{\partial p_x^*}{\partial \theta} > 0.$$

Proof. The planner observes $D_x(p)$ and can identify the sticky and non-sticky shares and the corresponding demands and utilities as described above. To simplify notation, suppress the index x , since the argument is identical for each demand state. We provide a monotone algorithm for determining p^* by induction on index k . Start with

$p_0^* = \bar{p}$ and $d_0^* = \bar{d}$, and construct iteratively

$$p_{k+1}^* = \frac{(1 - \theta)u(\bar{d}) + \theta u(d_k^*)}{(1 - \theta)\bar{d} + \theta d_k^*}, \quad d_k^* = d(p_k^*).$$

For any $\theta > 0$, the sequence $\{p_k^*\}$ is monotone. At $k = 0$, $p_1^* = u(\bar{d})/\bar{d} > p_0^* = \bar{p}$, implying $d_1^* < d_0^*$. Because $d'(p) < 0$ and $u(d)/d$ decreases in d , it follows that $p_{k+1}^* > p_k^*$ and $d_{k+1}^* < d_k^*$. The sequence converges for any $\theta < 1$, and the comparative statics follow directly from the limit expression for p^* . The second part of the proposition follows from $S_y(p)$ being strictly increasing in y . \square

Figure 1 illustrates an optimal rationing situation for two supply states, $y' < y$, and one demand state. The figure is drawn so that equilibrium F corresponds to the reference equilibrium: sticky and non-sticky demands both equal \bar{d}_x at the reference price \bar{p}_x . In this equilibrium, no rationing is required—condition (2.6) holds with strict inequality on the left, as aggregate utility exceeds expenditures.




Figure 1: An illustration of the optimal price-control and rationing policy for two supply states. Under the optimal price cap p_x^* , triangles ABC and CEF have equal areas, the latter weighted by the share of sticky consumers $(1 - \theta)$. Rationing reduces demand at p_x^* from Z to Y.

Now consider a negative supply shock, from y to y' , which moves the equilibrium from F to X . Is it optimal to allow the price to rise to this new level? Under the

price cap p_x^* , the demand of non-sticky consumers adjusts to point C . Aggregate utility U_z equals the area under AC , weighted by θ , plus the area under AF , weighted by $(1 - \theta)$. Expenditures equal the corresponding rectangles, yielding the difference between the right- and left-hand sides of condition (2.6) as the triangle CEF weighted by $(1 - \theta)$ minus the triangle ABC . Graphically, this difference increases with the price cap: it is negative at F and positive at A . Hence, there exists a unique p_x^* such that $\Delta ABC = (1 - \theta)\Delta CEF$. If, without rationing, $\mu_z = 1$, the equilibrium price is below p_x^* and no rationing is needed. If the unregulated price exceeds p_x^* , as in Figure 1, optimal rationing reduces demand to ensure that supply meets demand at the capped price.

There is therefore a unique p_x^* (independent of μ_z) such that the left inequality in (2.6) binds. This condition is fully determined by the demand schedule. Given p_x^* , the corresponding optimal rationing rate μ_z^* is set so that demand equals supply, $D_z = S_y(p_x^*)$. The same logic applies for any realization of demand; the price cap is state-contingent in x .

2.4 Parametric Demands

To illustrate the mechanism in a tractable form, consider a class of semi-log demand functions, to be estimated in the empirical application. We suppress the state notation (z, x, y) when not essential.

Proposition 2.3 (Semi-log case). *For the semi-log demand function*

$$d(p) = \exp\left(\frac{\alpha - p}{\beta}\right), \quad (2.8)$$

with non-negative parameters α, β , the optimal rationing price satisfies

$$p^* = \lambda(\theta)\beta + \bar{p}, \quad (2.9)$$

where $\lambda(\theta)$ solves

$$(\lambda - 1)e^\lambda = \frac{\theta}{1 - \theta}. \quad (2.10)$$

Proof. The semi-log specification implies the inverse demand $p = \alpha - \beta \ln d$, with

elasticity p/β . Utility follows as $u = \beta d + pd$, yielding indirect utility

$$u = v(p) = (\beta + p) \exp\left(\frac{\alpha - p}{\beta}\right).$$

Following the iterative procedure from Proposition 2.2, the updating rule for the optimal price cap is

$$p_{k+1}^* = \frac{(1 - \theta)(\beta + \bar{p}) \exp\left(\frac{\alpha - \bar{p}}{\beta}\right) + \theta(\beta + p_k^*) \exp\left(\frac{\alpha - p_k^*}{\beta}\right)}{(1 - \theta) \exp\left(\frac{\alpha - \bar{p}}{\beta}\right) + \theta \exp\left(\frac{\alpha - p_k^*}{\beta}\right)}. \quad (2.11)$$

Defining $\lambda_k = (p_k^* - \bar{p})/\beta$, this recursion becomes

$$\lambda_{k+1} = 1 + \frac{\theta \lambda_k e^{-\lambda_k}}{(1 - \theta) + \theta e^{-\lambda_k}}. \quad (2.12)$$

Starting with $\lambda_0 = 1$, the sequence $\{\lambda_k\}$ is strictly increasing and converges to the unique fixed point $\lambda(\theta)$ satisfying (2.10). \square

The optimal rationing price in (2.9) depends only on the expected price \bar{p} , the semi-log slope parameter β , and the share of non-sticky consumers θ . In empirical implementation, θ is identified jointly with α and β from estimated demand schedules. Proposition 2.3 shows directly how θ shifts the price cap: $\lambda(0) = 1$, $\lambda'(0) = 1/e \approx 0.368$, $\lambda(0.05) = 1.019$, $\lambda(0.1) = 1.039$, and $\lambda(0.5) = 1.28$. Hence, a 10% share of responsive consumers raises the distance between the optimal price cap and the reference price \bar{p} by roughly 4%. For any demand state x , the price cap increases one-to-one with \bar{p} and decreases with higher elasticity (lower β).

Finally, the policy can be expressed equivalently in terms of demand:

$$d^* = e^{-\lambda(\theta)} \bar{d}. \quad (2.13)$$

Thus, for this class of demand functions, the relative reduction of non-sticky demand before the price cap binds depends only on the share of responsive consumers. Non-sticky demand falls by at least 63% before rationing becomes active.

Consider next iso-elastic demand with $\sigma = p d'(p)/d(p)$.

Proposition 2.4 (Isoelastic case). *For a constant elasticity of individual demand*

$\sigma > 1$, let the aggregate elasticity of demand be $\varepsilon_D = \sigma s$, where $s = d/D_x$ denotes the share of responsive consumption in total demand. At the optimal p^* , it holds that

$$\frac{(1 - \theta) \left(\frac{p^*}{\bar{p}}\right) + \theta \left(\frac{p^*}{\bar{p}}\right)^{1-\sigma}}{1 - \theta + \theta \left(\frac{p^*}{\bar{p}}\right)^{1-\sigma}} = \frac{\varepsilon_D^*}{\varepsilon_D^* - s^*}, \quad (2.14)$$

where ε_D^* and s^* are evaluated at p^* . Moreover, both p^* and ε_D^* increase with θ .

We omit the proof – the result follows after a few steps from $p^* = U_z/D_z$ in Proposition 2.1.

The same observed ε_D can arise from different combinations of σ and θ , and therefore the policy depends on the decomposition between intrinsic elasticity σ and the share of responsive consumers θ . A higher θ raises the threshold price for intervention.

The price cap is minimal when all demand is sticky ($\theta = 0$), the result implies in that case that prices can increase by a factor $\sigma/(\sigma - 1)$ before rationing becomes optimal:

$$p^* = \frac{\sigma}{\sigma - 1} \bar{p}.$$

The associated reduction in non-sticky demand is

$$d^* = \left(\frac{\sigma - 1}{\sigma}\right)^\sigma \bar{d},$$

so $d^*/\bar{d} < 1/e$ and $d^*/\bar{d} \rightarrow 1/e \approx 0.368$ as $\sigma \rightarrow \infty$. The demand ratio d^*/\bar{d} increases with elasticity σ : non-sticky demand falls by at least 63% before the price cap binds, as in the case of semi-log demand.

The parametric demands show that the price cap policy remains well-defined even when all demand is inelastic ($\theta = 0$). To clarify the total welfare gain from the policy in this case and in general, it proves useful to state:⁷

Remark 2.1. For given demand D_x and supply S_y schedules and state z , the total increase in *ex post* surplus due to the policy is measured by XYZ in Fig 1.

⁷We omit the proof, which is not immediate; the complete argument appears in the working paper version of this paper.

2.5 Discussion and Roadmap

The main characterization establishes a simple rule: the planner regulates solely through a state-contingent rationing policy and a price cap. The cap p_x^* is pinned down by the average utility per unit of total demand; rationing μ_z^* is chosen to clear the market at p_x^* for the realized supply. Proposition 2.1 gives the complementarity condition: ration only when aggregate consumer surplus would turn negative. Proposition 2.2 shows how to implement the rule state by state, including comparative statics in the share of non-sticky consumers θ . The parametric illustrations (semi-log and iso-elastic) provide formulas for p_x^* .

The remainder of the theory section extends and applies this rule along three main margins. First, in Section 3 we relax the assumption of marginal-cost supply. As a baseline, we introduce the incentive payments for truthful supply first without a price cap (Section 3.1) and then with a cap (Section 3.2), and then introduce investments that materialize the supply (Section 3.3). Second, Section 4 extends the framework to endogenous selection into stickiness. Third, Section 5 incorporates distributional concerns and selection externalities.

3 Incentive Payments for Efficient Supply

When supply is competitive, the planner can set the price cap based solely on demand: the equilibrium price always reflects marginal cost. This assumption is unrealistic when suppliers hold market power, leading reservation prices to deviate from true marginal costs. Such strategic behavior would not matter if marginal costs were observable and truthful supply bids could be enforced. But even in this case, firms may have private information about investments that determine their marginal costs.

The incentive payments considered next provide a solution to both problems.

3.1 Incentive Payments without a Price Cap ($\theta = 1$)

Consider n firms $i = 1, \dots, n$ with cost functions $C_i(\cdot)$. Each firm submits a non-decreasing supply schedule $\hat{s}_i : \mathbb{R}_+ \rightarrow \mathbb{R}_+$. Aggregate supply is $\hat{S}(p) = \sum_{i=1}^n \hat{s}_i(p)$, and the market clears at the price \hat{p} solving $\hat{S}(\hat{p}) = D(\hat{p})$. For each firm i , define the

counterfactual clearing price \hat{p}_{-i} from

$$\sum_{j \neq i} \hat{s}_j(\hat{p}_{-i}) = D(\hat{p}_{-i}),$$

and define the residual demand

$$D_i(p) = D(p) - \sum_{j \neq i} \hat{s}_j(p).$$

In addition to market revenue, firm i receives a side transfer,

$$T_i(\hat{p}, \hat{p}_{-i}) = \begin{cases} \int_{\hat{p}}^{\hat{p}_{-i}} D_i(p') dp', & \text{if } \hat{p}_{-i} > \hat{p}, \\ 0, & \text{otherwise,} \end{cases}$$

so its payoff is

$$\pi_i = \hat{p} \hat{s}_i(\hat{p}) - C_i(\hat{s}_i(\hat{p})) + T_i(\hat{p}, \hat{p}_{-i}).$$

Recall that each C_i is differentiable and strictly convex. In addition, we require that each \hat{s}_i is right-continuous and non-decreasing. Together with continuous demand, these conditions ensure existence and uniqueness of \hat{p} .

Proposition 3.1 (Incentive payments without a price cap). *Under the mechanism above:*

- (i) **Dominant-strategy truthfulness.** *Each firm's dominant strategy is the truthful schedule $s_i^*(p) = (C_i')^{-1}(p)$. At the realized clearing price \hat{p} ,*

$$\frac{d\pi_i}{ds_i(\hat{p})} = \hat{p} - C_i'(\hat{s}_i(\hat{p})),$$

so the best response satisfies $C_i'(\hat{s}_i(\hat{p})) = \hat{p}$, independently of others' bids.

- (ii) **Efficiency.** *With truthful bids, the allocation maximizes total welfare*

$$W = \int_0^Q D^{-1}(q) dq - \sum_{j=1}^n C_j(q_j), \quad Q = \sum_{j=1}^n q_j.$$

(iii) **Payoff equals marginal contribution.** Under truthful bidding,

$$\pi_i = W - W_{-i},$$

where W_{-i} is welfare when firm i is removed and the market re-clears at \hat{p}_{-i} .

Proof. See Appendix F □

The proof follows (relatively) standard arguments for Vickrey–Clarke–Groves (VCG) mechanisms. The transfers $T_i(\hat{p}, \hat{p}_{-i})$ top up firms’ market revenues so that each firm’s total payoff equals its marginal contribution to welfare. Under strictly increasing marginal costs ($C_i'' > 0$), these contributions are subadditive, $\sum_i (W - W_{-i}) < W$; hence total transfers remain below aggregate surplus. Intuitively, the loss from removing one firm is partly offset by re-optimization of remaining suppliers.

The strong efficiency properties hinge on private information residing on the supply side while the planner observes demand. If both sides held private information, total surplus from trade would generally be insufficient to finance the transfers (Krishna and Perry, 1998), leading to a multi-unit version of the Myerson–Satterthwaite Theorem (Myerson and Satterthwaite, 1983).

3.2 Incentive payments with a price cap ($\theta < 1$)

We now extend the mechanism to the case of sticky demand ($\theta < 1$). The cap p^* is designed under the assumption of truthful supply. We therefore focus on mechanisms that ensure this assumption holds: they induce truthful supply when p^* binds.

When $\theta < 1$ but the cap does not bind, Proposition 3.1 applies as such. Assume thus that the cap binds. When firm i is removed, rationing must increase and production falls to $Q - \Delta$, defined by

$$C'_{-i}(Q - \Delta) = p^*.$$

The associated welfare change can be decomposed into two parts:

$$A = \int_0^{Q-\Delta} (C'_{-i}(q) - C'(q)) dq, \quad B = \int_{Q-\Delta}^Q (p^* - C'(q)) dq.$$

Term A represents the pure cost savings generated by firm i . The firm receives the

full value of these savings at the constant price p^* . Term B captures the net surplus from firm i 's participation, that is, the additional output enabled at price p^* . This surplus is also fully compensated through market revenue. Therefore, when the cap binds, no additional transfer is required, and the optimal mechanism sets $T_i = 0$.⁸

Proposition 3.2 (Incentive payments with $\theta < 1$). *Suppose p^* is the welfare-maximizing price cap. Under the Groves–Clarke pivot rule, if p^* binds:*

- (i) *Truthful supply is a dominant strategy for each firm,*
- (ii) *Each firm's payoff equals its marginal contribution to welfare*

$$\pi_i = W - W_{-i} = A + B,$$

- (iii) $T_i = 0$.

If p^ does not bind, Proposition 3.1 applies.*

The incentive payments defined here can therefore be implemented in the same way as under Proposition 3.1, but when the cap binds, the required transfer is zero.

The mechanism admits two implementation properties.

(i) Simplicity. The mechanism relies exclusively on repeated applications of the standard market-clearing rule. The planner computes the equilibrium price once with all firms and n additional times, each time removing a single firm to obtain the counterfactual equilibrium. These counterfactuals determine transfers but do not affect the final allocation or price.

(ii) Transfer reduction through price caps. In tight markets, counterfactual prices may rise substantially, implying large transfers in the absence of price regulation.

⁸If the cap does *not* bind with C but *does* bind with C_{-i} , we need to modify “ $A + B$ ” slightly. Let $P(q)$ be inverse demand. Define Q by $C'(Q) = P(Q) = p_0 < p^*$ (the cap does not bind), let Q' satisfy $P(Q') = p^*$ (so $Q' < Q$), and let $Q' - \Delta$ satisfy $C'_{-i}(Q' - \Delta) = p^*$ (so $\Delta > 0$). Rationing takes place over $[Q - \Delta, Q']$ and surplus change over $[Q', Q]$ is given by the Harberger triangle

$$B' \equiv \int_{Q'}^Q (P(q) - C'(q)) dq.$$

Thus, the total welfare change is $A + B + B'$.

A binding welfare-maximizing price cap limits this escalation and, as shown in Proposition 3.2, reduces transfers to zero. Firms are fully compensated for their marginal contributions through the capped market price. By this property, the mechanism is financially more disciplined than a system of incentive payments without the cap policy.

A large policy and regulatory literature, mostly in the electricity context (see Fabra, 2018), views price caps and capacity remuneration mechanisms as complementary instruments: caps limit exposure to high prices and market power, while capacity payments restore the scarcity rents needed for investment. In our framework, by contrast, the price cap p^* arises as the welfare-maximizing response to sticky demand under uniform pricing. This distinction is central for the investment stage: as we show next, the mechanism aligns firms' investment incentives with social welfare without relying on separate capacity payments. A broader discussion appears in the concluding section.

3.3 Investments on the Supply Side

The mechanism described above applies without an investment stage, but it is natural to ask how the price-cap policy affects producers' investment decisions that determine the observed supply schedules $S_y(p)$. We now introduce an investment stage preceding the realization of the market state $z = (x, y)$.

Each firm $i = 1, \dots, n$ chooses an investment level K_i at cost $G_i(K_i)$, where $G_i(\cdot)$ is convex and differentiable. Investment shifts the firm's operating cost function $C_i(q_i; K_i)$, which is convex in output and nonincreasing in K_i . The aggregate supply schedule is

$$S_y(p) = \sum_{i=1}^n s_{i,y}(p; K_i),$$

with each $s_{i,y}$ strictly increasing in price and in the supply state y .

Proposition 3.3 (Ex ante investment efficiency). *Suppose the planner implements the ex post rationing policy (p_x^*, μ_z^*) , defined in Propositions 2.1–2.2, and firm-specific incentive payments defined in Proposition 3.2. Under uniform pricing and non-discriminatory rationing, the policy (p_x^*, μ_z^*, T_i) maximizes expected welfare ex ante, subject to voluntary firm investment decisions with general convex costs $G_i(\cdot)$.*

Proof. Firm i 's profit in state z equals its marginal contribution to total welfare:

$$\pi_{i,z}(K) = W_z(K) - W_{z,-i}(K_{-i}).$$

Taking expectations over market states, firm i chooses its investment K_i to solve

$$\max_{K_i} \mathbb{E}_z[\pi_{i,z}(K)] - G_i(K_i).$$

The first-order condition for this problem coincides with the planner's maximization problem

$$\max_K \mathbb{E}_z[W_z(K)] - \sum_i G_i(K_i),$$

so equilibrium investment is efficient ex ante. In equilibrium, each firm invests until its private marginal benefit—its expected contribution to total welfare—equals its marginal investment cost. Since transfers ensure that firms internalize their marginal contributions, no other uniform-price mechanism can improve expected welfare. \square

Interpretations of investment. The investment variable K_i can be interpreted broadly as any firm-specific decision that shifts the cost function $C_i(q_i; K_i)$ or its curvature. In the simplest case, K_i represents capacity expansion that enlarges the feasible output set, as in traditional capacity investment. More generally, K_i may capture a technology choice, or an improvement in operational reliability that increases the expected availability of capacity while reducing its variance. It can also be viewed as a flexibility investment that flattens marginal costs and enables faster response to supply shocks, or as a hedging investment that reduces exposure to the carbon price by lowering the emissions intensity of production. Regardless of interpretation, all these decisions enter the market stage through the cost function $C_i(\cdot; K_i)$. Because the mechanism rewards each firm according to its marginal contribution to social welfare, $\pi_i = W - W_{-i}$, firms internalize the full social value of their investments. Hence, the same dominant-strategy property that guarantees efficiency ex post also ensures efficient investment ex ante, independent of the particular technological form of K_i .

Remark. If the market is perfectly competitive and firms are small (so investments are continuous and nonlumpy), marginal contributions are already internalized through the price mechanism, and efficient investment arises without additional transfers. This result parallels Makowski and Ostroy (1995), where agents rewarded according to

their marginal contributions implement efficient allocations. The present setting, with incentive payments designed by the planner, achieves this principle more generally: they provide each firm with its marginal contribution to social surplus, ensuring efficiency both ex post and ex ante.

A parallel issue arises on the demand side: consumers may choose whether to be price-responsive, and this choice is itself a welfare-relevant margin (e.g., Ito et al., 2023). In environments where the planner can discriminate among demand types, the efficient policy would ration only the sticky consumers—an insight going back to Wilson (1989) and Joskow and Tirole (2007). In our setting, however, uniform pricing rules out such discrimination. Given this constraint, the induced selection between responsive and sticky consumers is efficient under the ex post policy derived above. The next section establishes this result and after this we show how distortions arise once distributional objectives are introduced.

4 Endogenous Selection into Stickiness

The planner commits to the ex post policy characterized in Propositions 2.1–3.2 and applies it in every realized state $z = (x, y)$ using observed market schedules (D_x, S_y) to implement (p_z^*, μ_z^*) . Consumers choose whether to be price-responsive before z is realized and form rational expectations over the induced distribution of equilibrium prices p_z .

Let each consumer draw a private cost $c \geq 0$ from a continuous CDF H on \mathbb{R}_+ for observing and responding to price. This is the cost for becoming non-sticky. We can write $c(\theta) = H^{-1}(\theta)$ and assume both $H(c)$ and $c(\theta)$ to be strictly increasing differentiable functions over their domain. For state z , write the indirect (gross) surplus at price p_z as

$$\phi_z(p_z) \equiv \int_{p' \geq p_z} \mu_z d_x(p') dp',$$

well-defined under the assumptions stated. A responsive consumer attains $\phi_z(p_z)$ at the realized price; a sticky consumer consumes $\bar{d}_x = d_x(\bar{p}_x)$, where $\bar{p}_x \equiv \mathbb{E}[p_z | x]$, and attains $\phi_z(\bar{p}_x)$. Ex ante (before z is realized), the expected surplus gain from responsiveness is

$$\Delta(\theta) \equiv \mathbb{E}_z[\phi_z(p_z)] - \mathbb{E}_z[\phi_z(\bar{p}_x)],$$

where the expectation is over the joint distribution of (x,y) . We have assumed that all functions are smooth so that the surplus is also smoothly dependent on the share of responsive consumers. Moreover, we assume

$$\Delta(\theta) - c(\theta)$$

is strictly decreasing in θ .

Proposition 4.1 (Efficiency of selection). *Fix the optimal ex post policy rule (p_z^*, μ_z^*) , applied in every realized state z . Then the competitive equilibrium selection into responsiveness maximizes welfare achievable under uniform pricing.*

Proof. (i) *Ex post efficiency given θ .* Fix a share θ . By Propositions 2.1–3.2, the allocation induced by (p_z^*, μ_z^*) maximizes welfare in every realized state z , subject to uniform pricing and with truthful supply ensured. Thus, taking θ as given, the market implements the constrained-efficient ex post allocation.

(ii) *Planner’s ex ante selection margin.* By the envelope theorem (using quasi-linearity and the optimality of the ex post allocation for each θ), increasing the share of responsive consumers by a marginal amount changes expected welfare by the gross surplus gain of converting the marginal consumer from sticky to responsive:

$$\frac{d}{d\theta} \mathbb{E}_z[W_z(\theta)] = \mathbb{E}_z[\phi_z(p_z^*)] - \mathbb{E}_z[\phi_z(\bar{p}_x)] = \Delta(\theta).$$

The marginal cost of increasing the responsive share is $c(\theta) = H^{-1}(\theta)$. Hence the planner’s ex ante marginal condition is

$$\Delta(\theta) - c(\theta) = 0$$

whenever the optimum is interior. By assumption, $\Delta(\theta) - c(\theta)$ is strictly decreasing, so the optimum is unique (interior or not).

(iii) *Private choice and rational expectations.* The competitive selection condition is identical. Given rational expectations over the price distribution induced by the policy, a consumer with cost c becomes responsive if and only if $c \leq \Delta(\theta)$. Thus, in an interior equilibrium, the marginal consumer satisfies $c(\theta) = \Delta(\theta)$, while the same boundary conventions apply when the inequality is strict for all θ . Therefore the competitive selection into responsiveness coincides with the planner’s ex ante

optimum. Since the equilibrium θ coincides with the planner’s optimum, and the ex post allocation is efficient for this θ by part (i), the market outcome maximizes welfare under uniform pricing. \square

This result formalizes the idea that stickiness need not be an anomaly in this setting: the planner applies the ex post policy state by state, consumers anticipate the induced price distribution and optimally choose their responsiveness, and the resulting composition θ of sticky and responsive demand is welfare-optimal ex ante under uniform pricing.

In an alternative interpretation, consumers differ in the scale of their consumption. If the cost c of observing and responding to prices is a given fixed cost, a larger consumer has a lower cost per unit of demand and therefore a stronger private incentive to become responsive. Selection into responsiveness may therefore be systematically related to consumer size: high-consumption consumers, or consumers with better technologies for monitoring and adjustment, are more likely to become responsive, while smaller or less flexible consumers may remain sticky. To the extent that these differences correlate with income or wealth, selection into stickiness is also distributionally relevant. This observation motivates the next extension.

5 Inequality-aware welfare objective

We now extend the planner’s objective to account for heterogeneity in the marginal value of money. The idea is that the planner may be more concerned about the welfare of certain consumer types, as in Dworczak et al. (2021), Akbarpour et al. (2024), Pai and Strack (2022), and Ahlvik et al. (2024). Each individual has a welfare weight ω , which reflects the planner’s valuation of a unit of monetary surplus accruing to that individual. We assume that ω is unobservable at the individual level but that the planner knows its distribution. In particular, the planner observes the average welfare weights $\bar{\omega}_s$ and $\bar{\omega}_r$ for sticky and non-sticky consumers, respectively, and uses these weights to evaluate utility in a quasi-linear setting.

Based on the analysis of strategic supply, it is without loss of generality to describe the price-cap policy assuming truthful supply. We start by assuming that the investments have already been made and thus technology choices are given. The

aggregate utility is modified as:

$$U_z^\omega = \mu_z \left((1 - \theta) \bar{\omega}_s u_x(\bar{d}_x) + \theta \bar{\omega}_r u_x(d_x) \right), \quad (5.1)$$

and the planner's welfare objective becomes:

$$W_z^\omega = U_z^\omega - C_y(D_z). \quad (5.2)$$

This leads to the following extension of the optimal policy condition (proof omitted for brevity):

Proposition 5.1 (Inequality-aware price cap). *Under the inequality-weighted welfare objective, the optimal policy satisfies:*

$$p_z D_z(p_z) \leq U_z^\omega \perp \mu_z^* \leq 1. \quad (5.3)$$

If the price cap binds, the optimal price satisfies:

$$p_z^* = p_x^{*,\omega} = \frac{(1 - \theta) \bar{\omega}_s u_x(\bar{d}_x) + \theta \bar{\omega}_r u_x(d_x^*)}{(1 - \theta) \bar{d}_x + \theta d_x^*}, \quad (5.4)$$

where $d_x^* = d_x(p_x^{*,\omega})$.

The optimal price cap $p_x^{*,\omega}$ equates the welfare-weighted average utility per unit of demand across consumer types. If $\bar{\omega}_s > \bar{\omega}_r$ —that is, if the planner places greater weight on sticky consumers who are less price responsive and more vulnerable—then the optimal price cap is lower than in the unweighted case.⁹ This reflects the increased importance of shielding inelastic consumers from high prices when the planner values distributional equity.

However, once investments that determine responsiveness are endogenous, the distributive objective alters the nature of the optimal policy. In contrast to the main setting, the ex post price-cap mechanism may no longer coincide with the ex ante optimal mechanism. As shown, for instance, in Laffont and Tirole (1993) and Pavan et al. (2014), when early-stage decisions affect later allocations, the planner may wish to distort continuation policies to influence investment incentives. In Appendix F we

⁹Formally, this can be shown by using our fixed-point iteration $p_{k+1} = \Phi_\kappa(p_k) \equiv \frac{u(\bar{d}) + \kappa u(d(p_k))}{\bar{d} + \kappa d(p_k)}$, in which $\kappa \equiv \frac{\bar{\omega}_r}{\bar{\omega}_s} \cdot \frac{\theta}{1 - \theta}$, and observing that the fixed point is strictly increasing in κ .

show that, in the natural benchmark where the planner applies the same welfare weight to the self-selection cost and to surplus from market interaction, no such distortion is optimal: the optimal price-cap rule is unchanged.

6 Application: Energy crisis, France 2022–2023

6.1 Background

The European energy crisis emerged following the onset of the war in Ukraine, which disrupted supplies of gas, oil, and electricity and led to widespread economic impacts due to a surge in energy prices. Figure 2 illustrates the timeline and evolving expectations of the crisis’s severity, as reflected in France’s electricity prices. It shows the prices at which electricity for 2023 delivery were contracted in 2021–2022. Contract prices rose to more than ten times (!) their usual levels prior to 2022, imposing insurmountable cost burdens on entities needing to procure electricity in advance, such as those serving final consumers. The crisis threatened the stability of Europe’s integrated electricity market by endangering the solvency of firms with contractual commitments, with potential collateral demands exceeding one trillion euros.¹⁰ To avert a “Lehman Brothers scenario in the energy sector,” governments committed tens of billions in loans and guarantees.¹¹

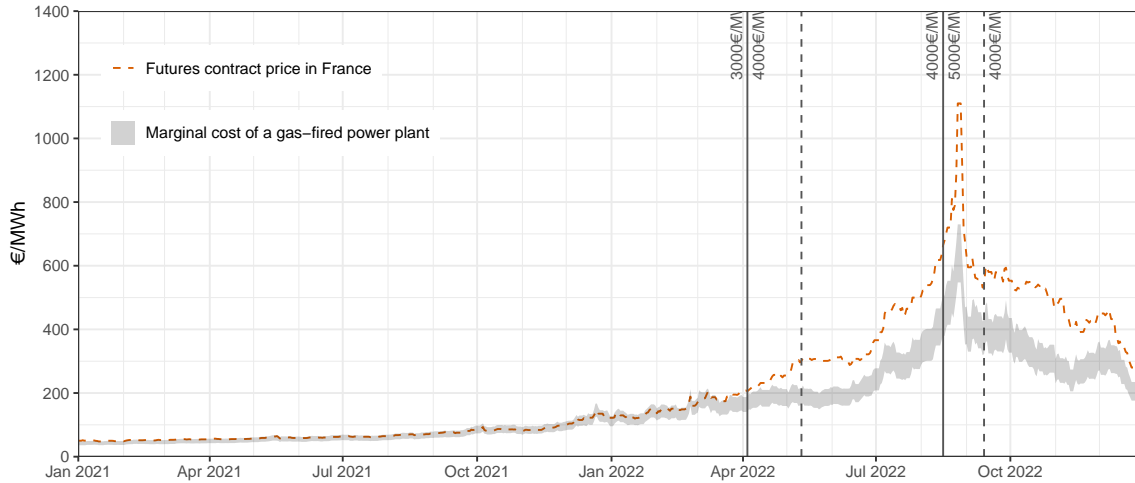
Three factors could plausibly account for these high expected price levels, alone or in combination. First, the war in Ukraine raised expected fossil input costs, especially natural gas, thereby increasing expected costs of thermal generation; Figure 2 visually depicts this factor as a gray band. Second, France’s nuclear fleet encountered a potentially systemic stress corrosion problem (corrosion-induced cracking), taking out about 40% of nuclear capacity at that same time and creating uncertainty about its near term availability (CRE, 2022a). Third, market participants may have anticipated scarcity episodes in which the market clearing price would be determined not by marginal costs but by the EU policy for setting and updating a price cap for wholesale electricity across the entire EU region. On April 4, 2022, the hourly wholesale price

¹⁰See *Bloomberg News*, “Energy Trading Stressed by Margin Calls of \$1.5 Trillion,” September 6, 2022. Available at <https://www.bloomberg.com/news/articles/2022-09-06/energy-trade-risks-collapsing-over-margin-calls-of-1-5-trillion>.

¹¹See Tom Wilson and Philip Stafford, “Why are Europe’s power producers running out of cash?,” *Financial Times*, September 6, 2022. Available at <https://www.ft.com/content/3a188669-7eeb-4154-91a8-f808ed8ced71>.

in France reached a threshold that triggered an automatic increase in the maximum clearing price under EU rules from €3,000/MWh to €4,000/MWh (CRE, 2022a). According to the protocol, the maximum price would be raised again under the same procedure if extreme wholesale price levels were to recur toward Winter 2022/23, a possibility that was salient in Summer 2022; the price cap might be raised repeatedly during the winter.¹²

Figure 2: Price expectations



Notes. Daily settlement prices for the 2023 French electricity futures contract, and the expected marginal cost range for gas fired generation based on unit efficiency and daily settlement prices for 2023 inputs (natural gas and EU ETS allowances). The range reflects typical combined-cycle efficiencies (45–60%). Source: own calculations. Vertical lines indicate EU wide revisions to the wholesale electricity price cap.

6.2 Crisis quantification

To gauge market expectations in August 2022, we construct counterfactual hourly price outcomes for 2023 under different combinations of the three factors and ask which of these outcomes are consistent with the 2023 futures prices observed in 2022. The prices and quantities that actually realized in 2023 come from the day-ahead market data, but we also reproduce them using hourly demand and supply schedules. For counterfactuals, we hold the hourly demand schedules as given by the data and

¹²Contemporary market commentary and survey evidence from the French regulator support this interpretation of market expectations (CRE, 2022b).

construct the supply curve for each hour of the year from engineering estimates combined with the counterfactual scenario assumptions. The constructed supply curve allows us to analyze separately the effects of higher input prices and reduced nuclear availability.

The data are drawn from the French day-ahead wholesale electricity market. For each hour between 1 January and 31 December 2023, we observe the complete set of demand and supply bids, price-quantity pairs, together with clearing prices and quantities. The dataset contains approximately 4.3 million anonymous bids, averaging about 500 per hour. We use unit-level information for 57 fossil-fuel plants obtained from RTE, the French transmission system operator, to build the supply estimates that depend on fossil inputs. These data report nameplate capacity (MW), technology, e.g., combined cycle gas turbine (CCGT) or steam turbine (ST), and commissioning year. We match technology and vintage to engineering-based heat-rate estimates to recover marginal cost curves. Appendix B provides details, but two features are noteworthy. First, marginal-cost estimates combined with observed demands reproduce the baseline 2023 hourly equilibrium once we include markups with magnitudes similar to those documented in Reguant (2014). Positive markups are not inconsistent with efficiency, as generators face start-up costs. Second, not all supply and demand clear through the exchange. To obtain an aggregate representation, we adjust the bid schedules using day-ahead forecasts of both total demand and supply; this procedure replicates the equilibrium outcomes up to rounding error.

The observed prices in 2023 yield a mean of €96/MWh. Table 1 contrasts the 2023 baseline with our quantification of how the three factors discussed above may have shaped crisis expectations. The first row reports the effect of higher input prices, set to the forward contract input prices observed in August 2022. Under this counterfactual, the mean expected price rises to €174/MWh due to higher input prices.¹³ The columns report results for alternative price-cap levels, reflecting the range of outcomes that the EU’s administrative price-cap protocol could have generated by 2023. Across these caps ranging from 3,000 to 7,000 €/MWh, expected prices remain unchanged: the increase in input costs does not raise the market price sufficiently for the cap to bind in any hour. We conclude that higher gas prices alone, while raising expected prices by roughly 60%, are insufficient to account for the expectations depicted in Fig. 2.

¹³Relative to realized 2023 input prices, August 2022 forward-contract levels were highest for gas-fired units (roughly triple), while coal- and oil-fired units were less than 50% above realized levels.

Table 1: EU price cap: Counterfactual mean hourly prices in 2023

Nuclear availability	Price cap (€/MWh)					Share of hours rationing
	3,000	4,000	5,000	6,000	7,000	
100%	174	174	174	174	174	0
90%	215	219	222	226	230	0.004
80%	304	323	342	360	379	0.019
70%	486	540	595	650	705	0.055
60%	847	1,001	1,156	1,310	1,465	0.155

Notes. The table reports simulated 2023 market outcomes by nuclear availability (rows) and price caps (columns). The share of hours with rationing applies to all columns 2–6. Input prices are fixed at 25 August 2022 closing levels. Nuclear availability is relative to realized 2023 nuclear generation. All prices are in €/MWh.

As nuclear availability declines from 100% to 60%—the level observed in August 2022—expected market prices rise sharply. In this table, input prices are held constant at their August 2022 forward levels, so variation across rows reflects only changes in nuclear availability relative to actual 2023 production. At 90% availability, expected prices rise to €215–€230/MWh, and at 80% they move into the €304–€379/MWh range. Further reductions generate steep increases: at 70% availability, expected prices lie between €486 and €705/MWh, and at 60% availability they reach €847–€1,465/MWh. The key mechanism is increased gas utilization: as nuclear output falls, gas-fired units enter the supply mix in more hours, raising the average expected price and amplifying the impact of high gas input costs.¹⁴ At sufficiently low nuclear availability, the increase in gas hours also raises the incidence of rationing hours, making outcomes increasingly dependent on the price-cap level—as reflected in the variation across columns. This dependence grows as nuclear availability falls. Although in August 2022 expected forward prices briefly exceeded €1,000/MWh and nuclear availability was close to 60%, we take 70% availability and a €7,000/MWh cap as approximating prevailing August expectations for 2023, yielding average contract prices of €705/MWh and total procurement costs of €330 billion.

When maximum supply falls short of minimum demand, prices are no longer pinned down by marginal cost alone but also by the administrative ceiling. In Table 1, the expected prices therefore reflect both marginal unit costs and the frequency with

¹⁴Conversely, at 100% nuclear, the counterfactual mean price (€174/MWh) remains below the marginal cost of gas because cheap nuclear limits the deployment of gas-fired units.

which the cap binds, with the latter reported in the last column. Importantly, these binding-cap (and hence rationing) frequencies are implied by the crisis expectations embedded in market prices, together with underlying demand and supply fundamentals. Notably, a 5% share of binding-cap hours implies that rationing was anticipated as a regular event—about one hour per day on average, with rationed quantities well below 10% in all cases (Appendix E). In the next section, in which we compute the welfare-optimal rationing rule, the resulting rationing incidence and quantities remain in the same ballpark.

6.3 Optimal price cap

We apply the optimal price-control mechanism to the 2023 fundamentals. In this counterfactual equilibrium, the price cap is set at its optimal level rather than at the EU cap, yielding a counterfactual mean price for 2023. This mean price corresponds to the contract price that market participants would have expected if the optimal policy, as we have formulated it, had been in place.¹⁵ We implement the mechanism in two ways that yield essentially the same results. First, we compute the counterfactual equilibrium numerically using the raw demand bid data and our counterfactual supply estimates. Second, instead of using the raw step-function bid schedules, we fit a parametric demand curve to each bid schedule and run the mechanism using the fitted curves. The parametric approach links the empirical analysis more directly to the theory and clarifies the quantitative determinants of the optimal cap.

For each hour, we observe a demand schedule as a step function constructed from a set of bid price-quantity pairs, denoted $(p_i, D_{xi})_i$, where i indexes points along the demand curve for any given hour x . We approximate this schedule using the parametric form:

$$D_{xi} = \gamma_x + \eta_x \exp\left(-\frac{p_i}{\beta_x}\right) + \varepsilon_{xi}, \quad (6.1)$$

where γ_x , η_x , and β_x are parameters to be estimated for each hour x , and ε_{xi} is an error term. Parameter γ_x [GW] measures the volume of sticky demand; parameter η_x [GW] measures responsive demand at price zero; and parameter β_x [€/MWh] measures the price at which responsive demand is reduced by about $1 - e^{-1} \approx 63\%$. We fit the

¹⁵Here, as in Fig. 1, contract is a futures contract on a fixed quantity at a fixed price.

formula in (6.1) to a curve, as the bids define a schedule, using maximum likelihood estimation. Dropping subscripts x, i and taking the estimated γ, η, β , and the reference price \bar{p} , the expected volume from (6.1) is $\bar{d} = \gamma + \eta \exp(-\bar{p}/\beta)$, which gives the sticky volume for $\bar{p} \rightarrow \infty$ and thereby the sticky share follows:

$$1 - \theta = \frac{\gamma}{\gamma + \eta \exp(-\bar{p}/\beta)}. \quad (6.2)$$

By defining $\alpha = \beta \ln(\eta/\theta)$ we can rewrite an estimate for the demand in terms of θ, α, β :

$$D = (1 - \theta)\bar{d} + \theta \exp\left(\frac{\alpha - p}{\beta}\right). \quad (6.3)$$

This is now a demand schedule with a sticky part and a non-sticky price-responsive part, with $d(p)$ taking a semi-log form as in (2.8) of illustration 2 in Section 2.4. The full set of steps for the procedure is:

1. Estimate demand (6.1) for each hour.
2. Impute the reference price \bar{p} .¹⁶
3. Compute the optimal price cap p^* as in Proposition 2.3 using the estimated demand parameters.
4. If p^* binds:
 - (i) activate the supply bids with reservation prices (weakly) below p^* , and
 - (ii) eliminate demand bids (pro-rata) until the demand volume equals the supply volume at p^* .

The key features of the estimated demand schedules can be summarized briefly. The mean estimated parameters are $\gamma \approx 46.0$ GW, $\eta \approx 8.0$ GW, and $\beta \approx 140.4$ €/MWh in 2023. The parametric form fits the bid schedules closely: replicating market outcomes using the estimated demand schedule, instead of the actual bids, leaves hourly equilibria nearly unchanged.

¹⁶The reference price anchors the sticky demand, reflecting consumption commitments made in the past: we construct the reference price as a rolling one-year average of realized market prices.

The elasticity of total demand takes a convenient form,

$$\varepsilon_D(p) = \frac{\eta p}{\beta(\eta + \gamma e^{p/\beta})},$$

which simplifies to $\varepsilon_D = p/\beta$ for the non-sticky component ($\gamma = 0$). To illustrate the ballpark magnitudes, we compute the optimal price cap using the mean estimated demand parameters ($\gamma \approx 46.0$, $\eta \approx 8.0$, and $\beta \approx 140.4$) and the average non-sticky demand share ($\theta \approx 0.026$). Substituting this value of θ into the defining equation $(\lambda - 1)e^\lambda = \theta/(1 - \theta)$ yields $\lambda(\theta) \approx 1.01$, which implies $\lambda(\theta)\beta \approx 142$ €/MWh. By Proposition 2.3, the corresponding optimal cap is therefore approximately 142 €/MWh above the reference price \bar{p} . This calculation is purely illustrative: in the empirical analysis we compute p^* separately for each hour of 2023 using the hour-specific demands, with details reported in the Appendix.

Table 2: Optimal price cap, mean prices, and rationing share

Nuclear availability	Optimal price cap (€/MWh)	Market price (€/MWh)	Share of hours rationing
100%	584	172	0.010
90%	604	202	0.029
80%	629	249	0.082
70%	661	334	0.201
60%	722	446	0.388

Notes. The table reports the optimal price cap, the resulting mean market price, and the share of hours in which rationing occurs in simulations for 2023. Input prices are fixed at 25 August 2022 closing levels. Nuclear availability is expressed relative to realized nuclear generation in 2023. All prices are in €/MWh.

Table 2 presents the expected 2023 market prices under the optimal cap as supply conditions tighten, mirroring the structure of Table 1. The cap in each hour is determined from the same demand schedules in each scenario, but the reference price varies across scenarios: as capacity becomes scarcer, the reference price rises, shifting up the implied cap. The number of hours in which the cap binds increases with the nuclear shortfall and is now well above the levels in Table 1, exceeding 20 % of hours under 70% availability. However, the quantities rationed in binding-cap hours still remain on average close to 10 percent of total load (Appendix E). Intuitively, when the supply curve is steeply rising, a relatively small reduction in demand lowers the

price to the desired level. The difference in the rationed amount required to bring the price down by hundreds is only a few percentage points.

6.4 Expenditures

We measure aggregate wholesale procurement expenditures and decompose them into (i) production costs, evaluated at estimated marginal-cost schedules, and (ii) payments in excess of marginal costs, consisting of scarcity rents (price–cost margins) and the incentive transfers required to sustain truthful supply. The incentive transfers are computed hour by hour from observed demand and estimated marginal costs, by evaluating each marginal unit’s contribution to total surplus; operationally, a unit’s marginal contribution equals its market revenue at the clearing price plus any additional payment needed to align its compensation with the social value it generates.

Table 3 shows how the price-cap mechanism affects the wholesale cost of procuring electricity. In the central case (70% availability), the optimal price-cap rule lowers expenditures by about two-thirds relative to a fixed €7,000/MWh cap—roughly €200 billion. For the same scenario, the welfare gain, calculated based on the area XYZ in Fig. 1, is €10.7 billion.

Table 3: Total expenditures and welfare (billion €)

Nuclear availability	Price cap rule		Expenditure difference	Welfare difference
	Fixed at €7,000/MWh	Optimal		
100%	72	71	1	0.0
90%	101	83	18	0.7
80%	180	103	77	3.5
70%	334	135	199	10.7
60%	648	172	476	29.7

Notes. Expenditures are the wholesale cost of procuring total electricity consumption in France at market prices for the hours for which we have full data ($N = 7,488$), scaled to a full-year equivalent ($N = 8,760$). The difference column reports the savings from optimal rationing relative to the fixed price-cap rule at €7,000/MWh. Welfare difference is calculated as the area XYZ in Fig. 1. Input prices are fixed at 25 August 2022 closing levels. Nuclear availability is expressed relative to realized nuclear generation in 2023. All values are in billions of euros.

Table 4 reports the incentive payments for the main marginal unit types in 2023 under two scenarios for nuclear availability (100 percent and 70 percent). For each

scenario, we compare a *baseline* case, in which the anticipated EU price cap of €7,000/MWh applies, with the *optimal* price-cap policy documented in Table 2. For each unit type, we also report the output and corresponding price–cost margin under the two policies.

Four patterns emerge. First, fossil-fuel generation is limited when nuclear output is high, but expands substantially when nuclear output is low. Second, the associated price–cost margins are low with abundant nuclear and rise sharply under nuclear shortfalls. Third, the optimal price cap compresses these margins by limiting prices in scarcity hours. Fourth, incentive payments are small throughout and are further reduced under the optimal price cap.

Fossil-fuel marginal units earn large markups precisely in the hours when the price cap binds. The cap is effective in a significant number of hours in which these units are deployed, so their average markup grows almost proportionally with the cap level: a cap of 7,000 €/MWh implies an average markup above 2,000 €/MWh for gas. A lower (optimal) cap therefore reduces markups roughly one-for-one and, in addition, curtails the costliest supply when marginal costs exceed the average utility per unit of consumption, limiting cumulative rent capture. In terms of quantities, output for marginal units is strongly shaped by variation in cheap supply: high-cost units produce only a fraction of output when low-cost nuclear is abundant (Panel A) but must be deployed at scale in scarcity, an above-tenfold increase (Panel B).

The final finding is that incentive payments required to implement truthful bidding remain small in all cases. Even at 70% availability, the plant level payments remain below €10/MWh under the optimal policy—an order of magnitude smaller than the corresponding markups. This reflects that most compensation is provided through market revenues and that the presence of many inframarginal units limits the marginal value of any individual generator.

7 Conclusion

In the electricity application, the mechanism is an exchange-level allocation rule: it prevents the exchange from clearing trades that generate negative surplus. This should not be confused with physical curtailment of final consumers. The day-ahead wholesale market clears bids submitted by retailers, aggregators, large consumers, and other market participants. A buyer whose bid is rationed must cover the shortfall through

Table 4: Price–cost markups and incentive payments under baseline and optimal cap

Units	Output (TWh)		Markup (€/MWh)		Incentive payments (€/MWh)	
	Baseline	Optimal	Baseline	Optimal	Baseline	Optimal
<i>Panel A. Nuclear availability 100%</i>						
Coal	0.23	0.23	196	150	16	7
Gas, CCGT	0.20	0.06	11	17	11	1
Gas, ST	0.03	0.01	8	10	3	1
Oil	0.18	0.18	226	143	3	2
<i>Panel B. Nuclear availability 70%</i>						
Coal	2.28	2.27	1,483	262	16	7
Gas, CCGT	7.27	3.59	2,184	145	16	7
Gas, ST	1.33	0.53	2,669	136	7	6
Oil	3.47	3.44	1,227	216	6	3

Notes. “Baseline” refers to the anticipated EU price cap (€7,000/MWh). “Optimal” refers to the optimal rationing mechanism. Output denotes total output from all units. Markup is calculated as the difference between total revenues and total costs divided by the output. Incentive payments are calculated as the difference between the sums of marginal contribution and markup and cost divided by the output. Nuclear availability is expressed relative to realized nuclear generation in 2023. Input prices are fixed at 25 August 2022 closing levels.

its own portfolio of contractual and operational instruments, including bilateral procurement, balancing-market purchases, storage or backup supply, or retail contracts that induce downstream demand reductions. Opt-in and time-varying pricing programs illustrate how retailers can activate household demand response through such downstream contracts (Fowle et al., 2021; Ito et al., 2021; Blonz, 2022). The efficiency gain from the exchange-level rule is therefore transmitted through the procurement, balancing, and retail arrangements of the affected buyers.

Housing provides a useful analogy, although not a direct application of the mechanism. Households make significant location commitments through school choices, jobs, care networks, and neighborhood ties. At the same time, rents or house prices may be set by a small flexible margin of movers, investors, or short-term rental demand (e.g., Genesove, 2003). From this perspective, temporary rent stabilization need not be viewed only as redistribution; it may also protect committed housing demand during a transient price shock, before location commitments can adjust. As in any market with a controlled price, scarcity is then allocated through non-price mechanisms—queues, search, landlord discretion, lotteries, or other forms of rationing.

In transportation, peak-period travel demand is often constrained by work schedules,

school trips, and location choices: the welfare cost of congestion or pollution may therefore not be inferred from the price that clears the marginal trip alone. Policies often rely on non-price rationing instruments rather than unrestricted congestion pricing. License-plate restrictions, driving bans, and low-emission zones ration access directly, either to reduce pollution or to manage congestion (Davis, 2008; Gallego et al., 2013; Viard and Fu, 2015; Barahona et al., 2020; Wolff, 2014; Galdon-Sanchez et al., 2023; Basso et al., 2021; Montero et al., 2025). These policies and their motivations differ from our mechanism, but they illustrate a related problem: when much of demand is committed, an unconstrained scarcity price may not be an efficient or acceptable way to allocate access.

References

- AHLVIK, L., M. LISKI, AND M. MÄKIMATTILA (2024): “Pigovian Income Taxation,” *CESifo Working Paper*.
- AKBARPOUR, M., E. BUDISH, P. DWORCZAK, AND S. D. KOMINERS (2024): “An economic framework for vaccine prioritization,” *The Quarterly Journal of Economics*, 139, 359–417.
- ASHCRAFT, A., N. GÂRLEANU, AND L. H. PEDERSEN (2011): “Two Monetary Tools: Interest Rates and Haircuts,” *NBER Macroeconomics Annual*, 25, 143–180.
- AUSUBEL, L. M., P. CRAMTON, M. PYCIA, M. ROSTEK, AND M. WERETKA (2014): “Demand Reduction and Inefficiency in Multi-Unit Auctions,” *The Review of Economic Studies*, 81, 1366–1400.
- BARAHONA, N., F. A. GALLEGO, AND J.-P. MONTERO (2020): “Vintage-Specific Driving Restrictions,” *The Review of Economic Studies*, 87, 1646–1682.
- BASSO, L. J., J.-P. MONTERO, AND F. SEPÚLVEDA (2021): “A Practical Approach for Curbing Congestion and Air Pollution: Driving Restrictions with Toll and Vintage Exemptions,” *Transportation Research Part A: Policy and Practice*, 148, 330–352.
- BERENTSEN, A. AND C. MONNET (2008): “Monetary Policy in a Channel System,” *Journal of Monetary Economics*, 55, 1067–1080.

- BLONZ, J. A. (2022): “Making the Best of the Second-Best: Welfare Consequences of Time-Varying Electricity Prices,” *Journal of the Association of Environmental and Resource Economists*, 9, 1087–1126.
- BORENSTEIN, S. AND S. P. HOLLAND (2005): “On the Efficiency of Competitive Electricity Markets With Time-Invariant Retail Prices,” *RAND Journal of Economics*, 36, 469–493.
- BULOW, J. AND P. KLEMPERER (2012): “Regulated prices, rent seeking, and consumer surplus,” *Journal of Political Economy*, 120, 160–186.
- CRE (2022a): “Analysis and Lessons Learnt Relating to the Price Peak on 4 April 2022,” *French Energy Regulatory Commission*.
- (2022b): “Electricity futures prices for winter 2022-2023 and the year 2023,” *French Energy Regulatory Commission*.
- DAVIS, L. W. (2008): “The Effect of Driving Restrictions on Air Quality in Mexico City,” *Journal of Political Economy*, 116, 38–81.
- DE FIORE, F., M. HOEROVA, C. ROGERS, AND H. UHLIG (2018): “Money Markets, Collateral and Monetary Policy,” NBER Working Paper 25319, National Bureau of Economic Research.
- DONNA, J. D. AND J.-A. ESPIN-SANCHEZ (2023): “The Illiquidity of Water Markets,” Tech. Rep. SSRN id 2667654, SSRN, paper posted December 1, 2023.
- DWORCZAK, P., S. D. KOMINERS, AND M. AKBARPOUR (2021): “Redistribution through markets,” *Econometrica*, 89, 1665–1698.
- FABRA, N. (2018): “A primer on capacity mechanisms,” *Energy Economics*, 75, 323–335.
- FOWLIE, M., C. WOLFRAM, P. BAYLIS, C. A. SPURLOCK, A. TODD-BLICK, AND P. CAPPERS (2021): “Default Effects And Follow-On Behaviour: Evidence From An Electricity Pricing Program,” *The Review of Economic Studies*, 88, 2886–2934.
- GALDON-SANCHEZ, J. E., R. GIL, F. HOLUB, AND G. URIZ-UHARTE (2023): “Social Benefits and Private Costs of Driving Restriction Policies: The Impact of

- Madrid Central on Congestion, Pollution, and Consumer Spending,” *Journal of the European Economic Association*, 21, 1227–1267.
- GALLEGO, F., J.-P. MONTERO, AND C. SALAS (2013): “The Effect of Transport Policies on Car Use: Evidence from Latin American Cities,” *Journal of Public Economics*, 107, 47–62.
- GENESOVE, D. (2003): “The Nominal Rigidity of Apartment Rents,” *Journal of Urban Economics*, 54, 154–168.
- GERLAGH, R., M. LISKI, AND I. VEHVILÄINEN (2025): “Pricing in Crisis – Dataset,” Zenodo, DOI: <https://doi.org/10.5281/zenodo.17964625>.
- GOWRISANKARAN, G., S. S. REYNOLDS, AND M. SAMANO (2016): “Intermittency and the value of renewable energy,” *Journal of Political Economy*, 124, 1187–1234.
- HORTAÇSU, A. AND D. MCADAMS (2018): “Empirical Work on Auctions of Multiple Objects,” *Journal of Economic Literature*, 56, 157–84.
- ITO, K., T. IDA, AND M. TANAKA (2021): “Selection on Welfare Gains: Experimental Evidence from Electricity Plan Choice,” .
- (2023): “Selection on welfare gains: Experimental evidence from electricity plan choice,” *American Economic Review*, 113, 2937–2973.
- JOSKOW, P. AND J. TIROLE (2006): “Retail electricity competition,” *The RAND Journal of Economics*, 37, 799–815.
- (2007): “Reliability and competitive electricity markets,” *The RAND Journal of Economics*, 38, 60–84.
- KRISHNA, V. AND M. PERRY (1998): “Efficient Mechanism Design,” Working paper, revised April 1998.
- LAFFONT, J.-J. AND J. TIROLE (1993): *A Theory of Incentives in Procurement and Regulation*, Cambridge, MA: MIT Press.
- LISKI, M. AND I. VEHVILÄINEN (forthcoming): “Redistribution through Efficiency: Theory and Evidence from Three Electricity Markets,” *RAND Journal of Economics*.

- MAKOWSKI, L. AND J. M. OSTROY (1995): “Appropriation and Efficiency: A Revision of the First Theorem of Welfare Economics,” *The American Economic Review*, 85, 808–827.
- MIER, M. (2024): “European electricity prices in times of multiple crises,” *Available at SSRN 4936684*.
- MONTERO, J.-P. (2008): “A Simple Auction Mechanism for the Optimal Allocation of the Commons,” *American Economic Review*, 98, 496–518.
- MONTERO, J.-P., F. SEPÚLVEDA, AND L. J. BASSO (2025): “Pricing Congestion to Increase Traffic: The Case of Bogotá,” *Journal of the European Economic Association*.
- MYERSON, R. B. AND M. A. SATTERTHWAITE (1983): “Efficient Mechanisms for Bilateral Trading,” *Journal of Economic Theory*, 29, 265–281.
- PAI, M. AND P. STRACK (2022): “Taxing Externalities Without Hurting the Poor,” Tech. rep., Cowles Foundation for Research in Economics, Yale University.
- PAVAN, A., I. SEGAL, AND J. TOIKKA (2014): “Dynamic Mechanism Design: A Myersonian Approach,” *Econometrica*, 82, 601–653.
- REGUANT, M. (2014): “Complementary Bidding Mechanisms and Startup Costs in Electricity Markets,” *The Review of Economic Studies*, 81, 1708–1742.
- RYAN, N. AND A. SUDARSHAN (2022): “Rationing the Commons,” *Journal of Political Economy*, 130, 210–257.
- SAH, R. K. (1987): “Queues, Rations and Market: Comparisons of Outcomes for the Poor and the Rich,” *American Economic Review*, 77, 69–77.
- TOKARSKI, F., S. D. KOMINERS, M. AKBARPOUR, AND P. DWORCZAK (2023): “A market-design response to the European energy crisis,” *mimeo, Northwestern University*.
- VÄLIMÄKI, T. (2003): *Central Bank Tenders: Three Essays on Money Market Liquidity Auctions*, no. E:26 in Bank of Finland Studies, Helsinki: Bank of Finland.

- VIARD, V. B. AND S. FU (2015): “The Effect of Beijing’s Driving Restrictions on Pollution and Economic Activity,” *Journal of Public Economics*, 125, 98–115.
- VICKREY, W. (1961): “COUNTERSPECULATION, AUCTIONS, AND COMPETITIVE SEALED TENDERS,” *The Journal of Finance*, 16, 8–37.
- WEITZMAN, M. L. (1977): “Is the Price System or Rationing More Effective in Getting a Commodity to Those Who Need it Most?” *The Bell Journal of Economics*, 8, 517–524.
- WHITESELL, W. (2006): “Interest Rate Corridors and Reserves,” *Journal of Monetary Economics*, 53, 1177–1195.
- WIJKANDER, H. (1988): “Equity and Efficiency in Public Sector Pricing: A Case for Rationing?” *Journal of Public Economics*, 36, 373–392.
- WILSON, R. (1989): “Efficient and Competitive Rationing,” *Econometrica*, 57, 1–40.
- WOLFF, H. (2014): “Keep Your Clunker in the Suburb: Low-Emission Zones and Adoption of Green Vehicles,” *The Economic Journal*, 124, F481–F512.

Part I

Appendix

Table of Contents

A	Data and replication package	40
B	Day-ahead market	40
C	Parametric demand estimation	41
D	Marginal-cost supply	41
E	Counterfactual simulations	43
F	Supplementary online material	2

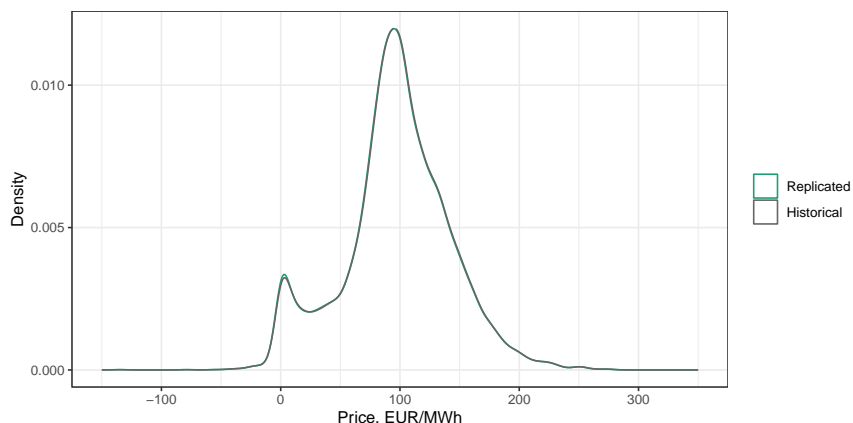
A Data and replication package

Available in Gerlagh et al. (2025)

B Day-ahead market

Market data. The quantification uses actual bids submitted to the French day-ahead market. The bid data is obtained from the European common electricity market’s Single Day-Ahead Coupling (SDAC) files, accessed through one of the common market operators, Nord Pool.

Figure B.1: Market price replication



Notes. Density plots of the historical hourly day-ahead spot prices in France in 2023 and the market prices as replicated by our model.

Replication. We construct price-responsive demand schedules from bid curves (price–quantity pairs) and replicate market outcomes using supply bids constructed in the same way. Replicating historical equilibrium outcomes serves as a robustness check on the data. We use the parsimonious linear program in Liski and Vehviläinen (forthcoming) to map submitted bids to equilibrium prices and quantities hour by hour over the sample period. The bid curves do not contain all information needed to fully replicate aggregate outcomes because they exclude imports and exports, non-standard bids, and transactions outside the exchange. We therefore apply level adjustments to demand and supply bid curves so that the implied equilibrium quantity matches the day-ahead load forecast—information available when bids are submitted (ENTSO-E

Transparency Platform). In counterfactual equilibria with rationing, we hold these adjustments fixed. This shortcut simplifies allocation under rationing; a regulator with access to the full algorithm could implement the mechanism EU-wide. The model reproduces historical market prices with high accuracy (Fig. B.1).

C Parametric demand estimation

We estimate the parametric form in Section 6.3, hour by hour, using each hour’s bid-curve data, i.e., the price (p_i) and quantity (D_{xi}) pairs. Table C.1 reports summary statistics for the MLE estimates of the parameter values, γ_x , η_x , and β_x , and for the imputed non-sticky demand share θ_x . To assess fit, we substitute the observed demand bid curve with the hour-specific parametric demand and recompute the market price for each hour, holding market supply as in data. The use of parametric demand yields clearing prices broadly consistent with historical outcomes: the mean of hourly prices using the parametric demand is €95.92/MWh cf. spot price of €96.65/MWh from the replication.

Table C.1: Parameter demand estimates

Statistic	N	Mean	St. Dev.	Min	Max
γ	8,760	45.971	10.051	23.425	80.531
η	8,760	8.049	4.767	0.112	26.350
β	8,760	140.400	159.663	12.593	2,605.899
θ	8,760	0.026	0.017	0.00000	0.115

Notes. The table reports mean, standard deviation, and the minimum and maximum values of the estimated parameter values γ , η , and β , and the imputed value of θ , for the N hours in the data.

D Marginal-cost supply

The counterfactual supply curve is constructed from plant-level data for large fossil-fuel units and aggregate production by energy source for other generation: hydro, nuclear, wind, solar, and bio/waste and other small units.

Marginal costs of fossil-fuel plants depend on fuel and carbon prices and heat rates (electric efficiency). Plant-level data are from RTE, the French TSO, and include start

date of operation, maximum capacity in megawatts (MW), and technology (combined cycle, steam turbine, and cogeneration status). RTE data do not report fuel sources. We use the Global Bioenergy Power Tracker (Sep 2025) and the Global Oil and Gas Plant Tracker (Jan 2025) to map plants with a capacity of over 20 MW to their fuel sources (combustion turbines can use coal or bio/waste; steam turbines can use natural gas or oil).

For input fuel and carbon prices, we use daily settlement prices from Refinitiv Eikon as follows: ICE Rotterdam Coal front-month futures for coal, EEX TTF Gas Day-ahead for natural gas, Brent FOB for oil, and EU ETS spot for carbon. Dollar values are converted to euros using the USD/EUR rate obtained through the European Central Bank API, and fuel energy content is converted to MWh using calorific values of 6,000 kcal per tonne of coal, 11.63 MWh per tonne of oil, and 7.37 barrels per tonne of oil. We obtain heat rates by mapping RTE technology and start date to the engineering values in Mier (2024), which provide typical heat rates by five-year construction cohort.


Supply costs include start-up and ramp-up costs and possibly other variable operation and maintenance costs. Firms may use block bids and other non-standard bid structures to recover the start-up costs, but their precise formulation is unobserved in the data. Start-up and ramping-up constraints generate an hourly pattern: nighttime and midday prices deviate downward, while morning and evening peak prices deviate upward, relative to daily marginal costs. We capture this pattern using the difference between a dynamic estimation that includes start-up costs and the competitive static estimate from Reguant (2014) to capture this pattern. To account for marginal costs beyond fuel and carbon, we scale the markups in Reguant (2014) by 1.5, increasing the mean bid-cost markup from .09 to .14.

The final fossil-fuel plant-level data set contains plant ID, nameplate capacity, fuel source and the associated emission factor, heat rate, and markup as defined above. Combined with the input price data, these variables yield plant-level reservation bids as quantity-price pairs. We account for outages via the ENTSO-E Transparency Platform, which reports unavailabilities for larger plants with capacities above 100 MW; smaller plants are assumed continuously available.

We recompute the 2023 market prices using observed demand bids and substituting market supply with cost-based reservation prices. These counterfactual prices align closely with the historical means and daily pattern, as shown in Fig. E.3. The main

differences are at the low end, which is unsurprising since we do not explicitly model the bidding by typically infra-marginal hydro, nuclear, and other assets. We validate the approach by comparing annual shares of marginal technologies from our simulation with the reported values by CRE, the French regulator: in 2023, gas fired technologies were marginal for 30% of hours in our simulation (30% for CRE), coal 10% (coal + borders 13%), and oil less than 1% ($< 1\%$) in 2023.¹⁷

Figure D.2: Historical market prices and marginal cost based prices



Notes. Comparison of 2023 French day-ahead spot prices: replicated market outcomes vs. marginal-cost-based supply. For each hour, the plot shows a box that represent the middle 50% of data, a line inside for the median, whiskers drawn at most 1.5 times the interquartile range from the box, and outliers beyond the whiskers as individual data points.

E Counterfactual simulations

The counterfactual simulations in the main text use observed market demand and the marginal-cost supply construction described above. We generate alternative counterfactuals by assuming: (i) input prices are fixed fixed at crisis-peak levels; (ii) nuclear availability is scaled back; and (iii) the price cap rule is varied.

Fixed input costs. The counterfactual simulations fix input prices for all days at their levels on 25 August 2022. For reference, the settlement price for a futures

¹⁷CRE, The monitoring and functioning of wholesale electricity and natural gas markets in 2023.

contract for electricity delivery in 2023 was €888/MWh (EEX), close to the mean closing price that week. On that date, input prices were €285.45/MWh for natural gas (TTF 2023 futures), \$368/t for coal, and \$102.58/bbl for oil; USD/EUR rate was .997. In the counterfactuals, these input prices are used to construct plant-level bids for the fossil-fuel fleet, as in the marginal-cost supply curve above.

Nuclear availability. The counterfactuals vary nuclear availability in 2023. For each scenario, we rerun the hourly market clearing by scaling the historical hourly nuclear output profile down by a fixed percentage across all hours. As a reference, when technical issues affecting much of the French nuclear fleet emerged in 2022, expected 2023 output was revised from 340–370 TWh to 300–330 TWh (February 2022).¹⁸ Realized nuclear output was 275 TWh in 2022 and 316 TWh in 2023.¹⁹

Fixed price cap. We run counterfactuals under two price-cap rules: a fixed, pre-determined cap and our optimal rationing rule. Under the fixed cap, we hold the cap constant throughout the year and vary its level; sticky demand is assumed to bid at the cap. We evaluate caps from €3,000/MWh to €7,000/MWh in €1,000/MWh increments. At the start of 2022, the price cap was €3,000/MWh with an automated rule: the cap increases by €1,000/MWh/MWh after five weeks if any hourly price in any market area exceeds 60% of the cap (=€1,800/MWh at the time).²⁰ This threshold was met on 4 April 2022, when France reached €2,987.78/MWh in hour 8–9, triggering a raise to €4,000/MWh on 11 May 2022.²¹ A second trigger occurred on 17 August 2022 when Baltic prices hit the EU ceiling, lifting the cap EU-wide to €5,000/MWh.²² The last increase was revoked before entering into force.²³

Optimal price cap. The optimal rationing rule updates the price cap for each hour. We calculate the optimal price cap directly based on the observed demand bid data by minimizing the trade-off between utility and expenditures, $\Delta ABC = (1 - \theta)\Delta CEF$ in Fig. 1. Table E.2 shows the summary statistics of the counterfactual market prices and the optimal price cap.

¹⁸EDF press releases, 11 Feb 2022, 18 May 2022, 3 Nov 2022.

¹⁹Source: ENTSO-E Transparency Platform.

²⁰Regulation (EU) 2019/943 and Commission Regulation (EU) 2015/1222 (CACM).

²¹SDAC Communication Note, 11 Apr 2022.

²²SDAC Communication Note, 23 Aug 2022.

²³ACER Decision 01/2023.

Table E.2: Optimal price cap

Statistic	N	Mean	St. Dev.	Min	Max
Market price	7,488	172.24	92.24	0.00	780.75
Optimal price cap	7,488	583.84	215.66	217.02	2,840.00

Notes. The table reports mean, standard deviation, and the minimum and maximum values of the counterfactual market clearing price and the optimal price cap for the N hours in the data.

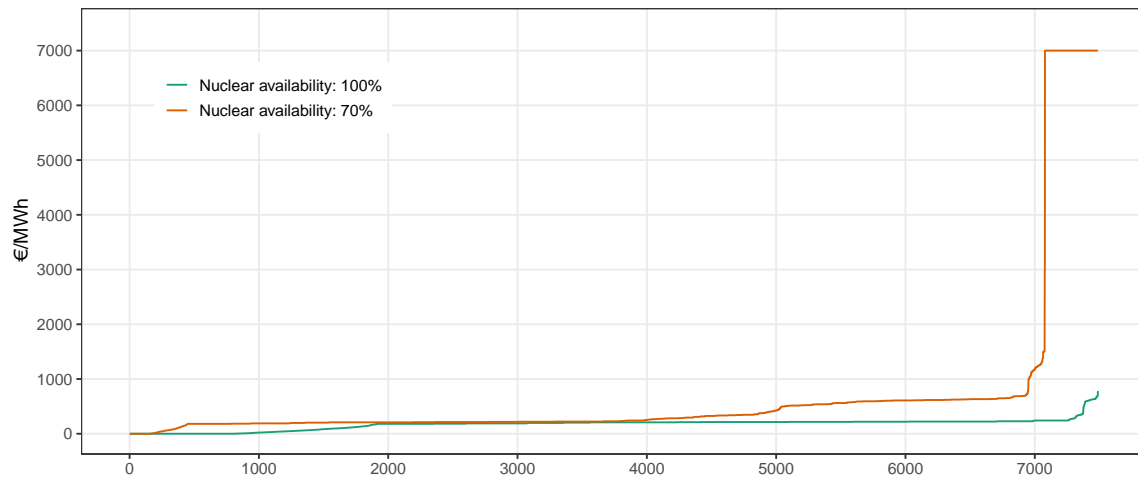
Table E.3 compares the share of quantity rationed with the fixed price cap rule and the optimal pricing rule as the nuclear supply reduces. In our central scenario on the crisis expectations, 70% of nuclear availability and €7,000/MWh fixed price cap, the mean rationed quantity, if rationing occurs, would have been 3,447 MWh. In the same scenario with our optimal price cap, the mean quantity rationed would have been 4,038 MWh.

Table E.3: Quantities rationed.

Nuclear availability	Price cap rule	
	Fixed at €7000/MWh	Optimal
100%	-	0.034 (0.001-0.076)
90%	0.017 (0.003-0.035)	0.051 (0.005-0.126)
80%	0.032 (0.002-0.078)	0.057 (0.004-0.152)
70%	0.051 (0.005-0.123)	0.071 (0.006-0.181)
60%	0.061 (0.003-0.155)	0.093 (0.008-0.217)

Notes. Table reports mean and 5%–95% interval of ratios between the quantity rationed to quantity demanded without rationing, calculated for each hour when rationing occurs, in different nuclear availability scenarios. Input costs are fixed to 25 August 2022 levels.

Figure E.3: Price distributions



Notes. Simulated hourly counterfactual prices for 2023 arranged from the lowest price hour in the left to the highest price hour on the right in two scenarios of nuclear availability.

Pricing in Crisis

Reyer Gerlagh et al., December 2025

SUPPLEMENTARY MATERIAL

F Supplementary online material

Proof of Proposition 3.1

Proof. We prove (i)–(iii) in order.

Step 1 (via a local variation at the clearing price). Fix opponents' bids \hat{s}_{-i} . Let \hat{p} be the clearing price under \hat{s}_i . Consider a small variation of \hat{s}_i that changes only the offered quantity at the realized clearing price, i.e., $\hat{s}_i^\varepsilon(p) = \hat{s}_i(p) + \varepsilon\phi(p)$ where ϕ is a smooth bump around $p = \hat{p}$ with $\int \phi = 1$. Denote by $\hat{p}(\varepsilon)$ the new clearing price and write derivatives at $\varepsilon = 0$.

Revenue term:

$$\frac{d}{d\varepsilon} [\hat{p}(\varepsilon) \hat{s}_i^\varepsilon(\hat{p}(\varepsilon))]_{\varepsilon=0} = \hat{s}_i(\hat{p}) \hat{p}'(0) + \hat{p} \frac{d\hat{s}_i(\hat{p})}{d\varepsilon}.$$

Cost term:

$$\frac{d}{d\varepsilon} C_i(\hat{s}_i^\varepsilon(\hat{p}(\varepsilon)))_{\varepsilon=0} = C_i'(\hat{s}_i(\hat{p})) \frac{d\hat{s}_i(\hat{p})}{d\varepsilon}.$$

Transfer term (note \hat{p}_{-i} does not depend on \hat{s}_i):

$$\frac{d}{d\varepsilon} T_i(\hat{p}(\varepsilon), \hat{p}_{-i})|_{\varepsilon=0} = - \left[D(\hat{p}) - \sum_{j \neq i} \hat{s}_j(\hat{p}) \right] \hat{p}'(0) = - \hat{s}_i(\hat{p}) \hat{p}'(0),$$

using market clearing $D(\hat{p}) = \hat{s}_i(\hat{p}) + \sum_{j \neq i} \hat{s}_j(\hat{p})$.

Summing the three derivatives, the $\hat{p}'(0)$ -terms *cancel*, yielding the Gateaux derivative

$$\frac{d\pi_i}{d\varepsilon} \Big|_{\varepsilon=0} = \left[\hat{p} - C_i'(\hat{s}_i(\hat{p})) \right] \frac{d\hat{s}_i(\hat{p})}{d\varepsilon}.$$

Thus, at the realized price \hat{p} , increasing quantity raises profit iff $\hat{p} > C_i'$, and lowers profit iff $\hat{p} < C_i'$. Hence the unique best response at \hat{p} satisfies

$$C_i'(\hat{s}_i(\hat{p})) = \hat{p},$$

independently of \hat{s}_{-i} . Because opponents' bids can make the clearing price equal to any p in the support, the only strategy that is optimal for every possible \hat{p} is the pointwise truthful schedule $C_i'(\hat{s}_i(p)) = p$ for all p . This proves (i).

Step 2 (Efficiency under truthful bidding). If each firm bids $s_i^*(p) = (C_i')^{-1}(p)$, then $\hat{S}(p) = \sum_i (C_i')^{-1}(p)$ is the aggregate marginal-cost schedule. The clearing condition

$\hat{S}(\hat{p}) = D(\hat{p})$ is the first-order condition for maximizing W : set total quantity Q so that $D^{-1}(Q) = \text{common marginal cost} = \hat{p}$, with quantities allocated across firms by $C'_i(q_i) = \hat{p}$. By strict convexity of costs and strict monotonicity of demand, the solution is unique and efficient. This proves (ii).

Step 3 (Payoff equals marginal contribution under truthful bidding). We establish the “area identity,” then conclude. Under truthful bids,

$$q_j = \hat{s}_j(\hat{p}) = (C'_j)^{-1}(\hat{p}), \quad q_j^{-i} = \hat{s}_j(\hat{p}_{-i}).$$

A change of variables $q = D(p)$ gives

$$\int_{Q_{-i}}^Q D^{-1}(q) dq = \int_{\hat{p}_{-i}}^{\hat{p}} p' D'(p') dp' = \left[p' D(p') \right]_{\hat{p}_{-i}}^{\hat{p}} - \int_{\hat{p}_{-i}}^{\hat{p}} D(p') dp'. \quad (\text{F.1})$$

For $j \neq i$, convexity and truthfulness imply

$$C_j(q_j) - C_j(q_j^{-i}) = \int_{q_j^{-i}}^{q_j} C'_j(q) dq = \int_{\hat{p}_{-i}}^{\hat{p}} p' d\hat{s}_j(p') = \left[p' \hat{s}_j(p') \right]_{\hat{p}_{-i}}^{\hat{p}} - \int_{\hat{p}_{-i}}^{\hat{p}} \hat{s}_j(p') dp'. \quad (\text{F.2})$$

Summing (F.2) over $j \neq i$ gives

$$\begin{aligned} \sum_{j \neq i} [C_j(q_j) - C_j(q_j^{-i})] &= \left[p' \sum_{j \neq i} \hat{s}_j(p') \right]_{\hat{p}_{-i}}^{\hat{p}} - \int_{\hat{p}_{-i}}^{\hat{p}} \sum_{j \neq i} \hat{s}_j(p') dp' \\ &= \hat{p} [D(\hat{p}) - \hat{s}_i(\hat{p})] - \hat{p}_{-i} D(\hat{p}_{-i}) - \int_{\hat{p}_{-i}}^{\hat{p}} \sum_{j \neq i} \hat{s}_j(p') dp'. \end{aligned} \quad (\text{F.3})$$

Subtracting (F.3) from (F.1) yields

$$\begin{aligned} \int_{Q_{-i}}^Q D^{-1}(q) dq - \sum_{j \neq i} [C_j(q_j) - C_j(q_j^{-i})] &= \left(\hat{p} D(\hat{p}) - \hat{p}_{-i} D(\hat{p}_{-i}) \right) - \left(\hat{p} [D(\hat{p}) - \hat{s}_i(\hat{p})] - \hat{p}_{-i} D(\hat{p}_{-i}) \right) \\ &\quad - \int_{\hat{p}_{-i}}^{\hat{p}} D(p') dp' + \int_{\hat{p}_{-i}}^{\hat{p}} \sum_{j \neq i} \hat{s}_j(p') dp'. \end{aligned}$$

The boundary terms simplify to $\hat{p} D(\hat{p}) - \hat{p} [D(\hat{p}) - \hat{s}_i(\hat{p})] = \hat{p} \hat{s}_i(\hat{p})$. The integral terms

combine as

$$-\int_{\hat{p}_{-i}}^{\hat{p}} D(p') dp' + \int_{\hat{p}_{-i}}^{\hat{p}} \sum_{j \neq i} \hat{s}_j(p') dp' = -\int_{\hat{p}_{-i}}^{\hat{p}} \left[D(p') - \sum_{j \neq i} \hat{s}_j(p') \right] dp' = \int_{\hat{p}}^{\hat{p}_{-i}} D_i(p') dp',$$

where $D_i(p') = D(p') - \sum_{j \neq i} \hat{s}_j(p')$ is the residual demand for i . Therefore,

$$\int_{Q_{-i}}^Q D^{-1}(q) dq - \sum_{j \neq i} [C_j(q_j) - C_j(q_j^{-i})] = \hat{p} \hat{s}_i(\hat{p}) + \int_{\hat{p}}^{\hat{p}_{-i}} D_i(p') dp'.$$

Subtracting $C_i(q_i) = C_i(\hat{s}_i(\hat{p}))$ from both sides gives

$$W - W_{-i} = \hat{p} \hat{s}_i(\hat{p}) - C_i(\hat{s}_i(\hat{p})) + \int_{\hat{p}}^{\hat{p}_{-i}} D_i(p') dp' = \pi_i,$$

as claimed. This proves (iii). \square

Proof of ex-ante optimality of price cap rule for heterogeneous welfare weights in Section 5.

Let the demand state x be continuously distributed with full support over a compact interval. Assume also that the price-cap rule p_x^* is a continuous and differentiable function. The planner applies this rule ex post in each state, but anticipates its ex ante impact on responsiveness and investment. The planner's problem is

$$\max_{\{p_x^*\}} \mathbb{E}_z[W_z^\omega] - \int_0^{c^*} \omega(c) c dH(c), \quad (\text{F.4})$$

where

$$W_z^\omega = \mu_z(p_x^*) \left[(1 - \theta) \bar{\omega}_s(\theta) u_x(\bar{d}_x) + \theta \bar{\omega}_r(\theta) u_x(d_x(p_z)) \right] - C_y(D_z), \quad (\text{F.5})$$

with conditional welfare weights

$$\begin{aligned} \bar{\omega}_r(\theta) &= \frac{1}{\theta} \int_0^{c^*} \omega(c) dH(c), \\ \bar{\omega}_s(\theta) &= \frac{1}{1 - \theta} \int_{c^*}^\infty \omega(c) dH(c), \end{aligned}$$

and with the selection margin satisfying $\theta = H(c^*)$.

After perturbing the optimal (differentiable) rule in a given state x , the first-order condition can be written as

$$\frac{\partial \mathbb{E}_z[W_z^\omega]}{\partial p_x^*} + \left(\frac{\partial W_z^\omega}{\partial \theta} - \omega(c^*) c^* \right) h(c^*) \frac{\partial \Delta}{\partial p_x^*} = 0, \quad (\text{F.6})$$

where

$$\frac{\partial c^*}{\partial p_x^*} = \frac{\partial \Delta}{\partial p_x^*}, \quad \frac{\partial \theta}{\partial p_x^*} = h(c^*) \cdot \frac{\partial c^*}{\partial p_x^*}.$$

To interpret the condition, note that (F.6) decomposes the planner's first-order condition into two channels (separated by the plus sign). The first measures the deviation from the ex-post optimal price cap (5.4). The second measures the deviation from optimal endogenous selection of responsiveness for uniform consumer weights (??). That is, the FOC informs us that if heterogeneous welfare weights change the optimal selection of responsiveness θ , the optimal price cap rule (5.4) must be adjusted as well. Yet as we will see, such is not the case. The optimal price cap rule is preserved.

The first term in brackets of (F.6) reflects the social benefits of adding the marginal consumer $c^* = \Delta$ from sticky to responsiveness. As we have seen in (??), the social value equals the private costs, denoted by the second term in brackets of (F.6). Heterogeneous weights does not change that, as both benefits and costs receive the same marginal consumer's weight. Thus the term in brackets is zero, and thus the ex-post optimal price cap rule is also preserved ex-ante.