Redistributive Peak Load Pricing

Leopold Monjoie

[1] Aalto University

[2] leopold.monjoie@aalto.fi

Abstract.

This article studies peak-load pricing in essential-goods markets such as electricity, transport, and network industries, when consumers have private information about their willingness to pay, belong to observable categories, and a market designer has redistributive objectives. I characterize the optimal mechanism and answer the following: who benefits from capacity expansion, and how redistributive preferences shape the optimal allocation. I derive structural conditions under which the consumers who gain or lose from the mechanism do not coincide with redistributive priorities. This occurs because allocating to one type propagates to others through informational rents and capacity scarcity, and these effects are evaluated at different social values. The same mechanism governs three policy results. First, I characterize the optimal tagging rule for peak-load pricing and show how observable categories should be ranked by capacity allocation and budget contribution. Second, in contrast to the utilitarian case, I establish a new peak-load investment rule in which the distortion from private information and redistribution persists at the optimum. Third, I further discuss the optimal nonlinear tariff that implements the mechanism, showing that spot pricing fails and that optimal block tariffs are generically neither increasing nor decreasing for consumers with strong redistributive priority.

1 Introduction

When demand is uncertain and capacity is limited, the standard framework of peak-load pricing provides a mechanism for determining who receives the scarce good and who finances the capacity needed to alleviate that scarcity. These questions are important for essential goods (in particular electricity, network industries, and health services), because inadequate rationing and capacity shortages in these markets impose costs that extend far beyond the consumers directly rationed.[1] The issue is particularly relevant because these industries are characterized by structural pressure to expand capacity, driven by demand growth, technological change, and the transition to low-carbon supply.[2] In that context, the rule that governs rationing today simultaneously determines the revenue stream that finances capacity, making the allocation of scarce supply and the financing of expansion part of the same economic problem. But these are also the sectors in which distributional concerns are strongest. When capacity is scarce, the allocation rule determines who is served; when capacity is expanded, the financing rule determines who bears the burden of paying for it. This creates a normative tension as the logic of peak-load pricing points toward allocating scarce capacity to the users with the highest marginal willingness to pay and financing new investment through a unique price that reflects aggregate scarcity. For many essential goods, however, policymakers are reluctant to let either access during shortages or the burden of financing expansion be determined solely by willingness to pay.[3] As a result, the objectives of efficiency, cost recovery, and redistribution need not align. The consumers who should receive the good under an efficient scarcity allocation are not necessarily those a redistributive regulator would want to prioritize, and those from whom investment costs are easiest to recover are not necessarily those on whom society would want the burden to fall.

To study this tension, I consider a market designer, for instance, a regulated firm acting under a mandate to serve consumers, who must choose capacity and a state-contingent allocation mechanism for a good with stochastic demand. The good is subject to an ex post hard capacity constraint[4] and the capacity must be financed under an ex ante budget constraint. Consumers belong to observable categories, but within each

[1]In electricity markets, (Hinchberger et al., 2024) find a deadweight loss of about \$2 billion annually from flat retail pricing over several US markets, while the 2021 Texas blackout caused \$80–130 billion in economic losses (DallasFed, 2021; Busby et al., 2021). The 2025 *Urban Mobility Report* reports annual congestion costs of \$269 billion in U.S. urban areas (Schrank et al., 2025). In vaccine allocation, Moore et al. (2022) showed that misallocation of vaccines during COVID-19 led to about 1.3 million direct deaths worldwide by the end of 2021.

[2]In electricity generation, the IEA projects almost 4,600 GW of additional renewable capacity worldwide between 2025 and 2030 (IEA, 2025). For electricity grids, the European Commission estimates total investment needs of around €1.2 trillion by 2040 (EuropeanCommission, 2025). In urban transport, the OECD estimates average investment needs of USD 400 per capita per year in high-income cities between 2015 and 2050 (Wagner, 2018).

[3]Redistributive intervention takes several forms in practice. In energy markets, governments have used price caps and subsidies to protect consumers from price spikes (Fetzer et al., 2024; Amores et al., 2025), while nonlinear tariffs have been designed to shift revenue burdens toward higher-usage consumers (Borenstein, 2012; Levinson and Silva, 2022; Smith, 2022; Randriamaro and Cook, 2026). In transport, the distributional impact of congestion pricing depends heavily on how revenues are recycled (Eliasson and Mattsson, 2006). In health emergencies, redistribution takes the form of non-price rationing schemes based on vulnerability or social role rather than willingness to pay (Wang et al., 2023; Akbarpour et al., 2023; Pathak et al., 2024).

[4]Capacity is determined before the state of the world is realized, but the allocation may depend on that realization. This captures the fact that infrastructure and financing decisions in markets such as electricity, transport, or health services must be taken before demand conditions are known, even though the actual use of scarce capacity can adjust once those conditions are observed. A hard capacity constraint means it is never optimal for the market designer to allow excess demand to materialize.

category, they privately know their willingness to consume, summarized by a single preference parameter that scales the value they derive from receiving a given quantity. The designer evaluates allocations using welfare weights placed on consumer surplus. These weights are treated as exogenous reduced-form objects rather than derived from explicit microfoundations. The interpretation is that the designer cannot directly observe whether a consumer is vulnerable or socially deserving and must instead rely on the willingness-to-pay information elicited through the mechanism as a proxy for social need.[5] The analysis proceeds in the following steps. I begin by taking the allocation of capacity and revenue requirements across categories as given and characterizing the optimal mechanism for a single category. This analysis identifies two sorting margins: one governing the incidence of capacity expansion, and one governing how redistributive preferences translate into reallocations at fixed capacity. These margins are summary measures of the constrained optimum. They first organize the within-category mechanism and then determine how capacity and revenue should be allocated across categories.

The main tension comes from the interaction between incentive constraints, redistributive preferences, and capacity scarcity. Conditional on capacity and revenue requirements, I first characterize the optimal mechanism for a single category and then study how the constrained optimum responds to changes in capacity and redistributive preferences. Beyond alleviating scarcity, capacity expansion affects the surplus generated by the mechanism, which can be redistributed to consumers through transfers. However, incentive compatibility requires the mechanism to leave rents to higher types, and once the designer assigns different social values to consumption and transfers across types, those rents can reduce the surplus available for redistribution while generating limited social value. This interaction yields the *capacity-sorting margin*. I show that the sign of this margin, together with the ordering of types by their ability to generate surplus, governs the set of winners and losers from a capacity expansion, and that this incidence depends on the level of capacity. To illustrate, suppose the designer gives higher priority to consumers with low demand. If these consumers are also more responsive to capacity expansion, relaxing scarcity may increase the informational rents that must be left to higher types. The surplus available for redistribution can then fall, so the consumers targeted by redistribution may lose from additional capacity.

I then compare the optimal allocations across redistributive priorities, holding capacity fixed. The comparison operates through two channels. First, redistributive priorities change the designer's effective social weights, which measure the social return from serving each consumer. Second, because capacity is fixed, assigning relatively more social value to some types necessarily reallocates scarce capacity away from others. The relevant object is therefore not the absolute difference in effective social weights, but their proportional difference across consumers. This yields the *preference-sorting margin*. I provide conditions

[5]In this article, higher welfare weights on higher types are equivalent to favoring consumers with a higher demand. For a related critique of using observed electricity demand as a proxy for consumers' social value, see Borenstein (2025), who shows that much of the variation in residential electricity use reflects household size, rooftop solar, and climate rather than imprudent consumption. Following Akbarpour et al. (2024), restricting attention to weights that depend only on the observable category i and the elicitable type θ is without loss. Since λ does not enter agents' preferences, no incentive-compatible mechanism can profitably condition allocations or transfers on a separately reported welfare characteristic.

under which average welfare weights and the direction of redistributive priorities determine the sign of this margin. I then show that, even when those priorities shift toward lower or higher types, the induced allocation need not move in the same direction. Reweighting one part of the type distribution propagates through incentive constraints to the rest of the allocation, so the consumers who receive greater social weight need not be the ones who benefit from the induced reallocation.

Using the same two sorting margins that organize the single-category mechanism, I then characterize the optimal allocation across categories: how the global capacity and revenue requirements are assigned across observable groups. Capacity is allocated according to the marginal social value of relaxing scarcity. I show that this marginal value is governed by the same *preference-sorting margin*, rather than by average welfare weights or absolute effective social weights. Revenue requirements are governed by a different logic. The tagging problem distinguishes between categories of contributors and non-contributors. A contributing category is assigned a positive revenue requirement so that its category budget constraint binds, and the mechanism generates no surplus to redistribute to its consumers. By contrast, the budget constraint of a non-contributing category is slack, and any surplus generated by the mechanism is redistributed within the category. I show that a category contributes only when the common social value of public funds at the equilibrium exceeds the redistributive cost of extracting revenue from that category, captured by its average social weight and by the category-level shadow value of the budget constraint. The extensive margin of revenue extraction, therefore, depends only on average welfare weights. The intensive margin, by contrast, depends on the interaction between the two sorting margins, since a category's ability to raise revenue reflects both how capacity is allocated to it and how its redistributive priorities sort consumption across types.

The final result establishes the incomplete-information counterpart of the canonical peak-load investment rule and shows how the optimal rule departs from its utilitarian benchmark. When consumers share the same welfare weight, informational rents are welfare-neutral, so the complete- and incomplete-information investment rules coincide at the optimum capacity. Under redistributive preferences, by contrast, the informational wedge persists even at the optimum. The new optimality condition equates the marginal investment cost with the incentive-adjusted social value of marginal consumption, aggregated across contributing categories and normalized by the equilibrium marginal cost of funds. The comparative statics of this investment rule are governed by the same two sorting margins developed above. Relative to the utilitarian benchmark, redistributive priorities raise optimal investment when the aggregate *preference-sorting margin*, which captures the marginal value of capacity, is positive and the aggregate *capacity-sorting margin* does not raise the equilibrium marginal cost of funds too much. The reverse holds when the opposite force dominates. The effect of a global capacity expansion on winners and losers within each category is then governed by the comparison between the category-specific *capacity-sorting margin* and the marginal investment cost.

I conclude by discussing a tariff representation of the mechanism to illustrate its policy implications and show how the same sorting margins that govern the optimal mechanism also shape nonlinear marginal price schedules within and across categories. This provides a policy interpretation of the results, especially for block tariffs commonly used in essential goods sectors to achieve redistribution. Redistribution not only implies that spot pricing fails to implement the optimum, even when combined with individualized lump-sum transfers, but also does not translate into uniformly lower prices for favored consumers. Instead, I show that it can give rise to non-standard block tariff structures, differentiated menus across observable groups, and ambiguous price and surplus responses to capacity expansions.

Related Literature The classic peak-load pricing literature studies the efficient allocation of a scarce, non-storable good across states when capacity is fixed in the short run and costly to expand in the long run (Boiteux, 1949; Steiner, 1957; Vickrey, 1963). Off peak, the optimal allocation sets users' marginal utility equal to operating marginal cost; on peak, capacity binds and the good is allocated so that marginal utilities are equalized at a common scarcity value, while optimal investment equates the marginal cost of capacity to that scarcity value (Crew et al., 1995; Vickrey, 1969).[6] This benchmark extends to incomplete information under a utilitarian objective. When consumers privately know their willingness to pay, the efficient allocation and investment can still be implemented through self-selection mechanisms (Spulber, 1992b; Chao and Wilson, 1987).[7] This article shows how the benchmark changes once the designer has redistributive objectives. In a more general mechanism-design environment, I characterize the optimal allocation when consumers differ both in their willingness to pay and in their social priority, and show that peak-load pricing no longer equates marginal utilities to a common scarcity value. Instead, marginal utilities are distorted by effective redistributive weights, so the utilitarian spot-pricing logic breaks down even when individual transfers are feasible.[8]

In the peak-load pricing literature, redistributive objectives under private information have received relatively little attention. Spulber (1992a) characterizes the self-selection contracts that implement the peak-load optimum under welfare weights, but does not study the comparative statics of redistribution, the incidence of capacity expansion, or the optimal investment rule.[9] More broadly, the article relates to the redistributive mechanism design literature, which studies similar tensions in settings such as externality

[6]One important part of this literature concerns implementation, which I illustrate in the discussion. The standard result is that the efficient utilitarian allocation can be decentralized through responsive or spot pricing. Prices equal operating marginal cost off-peak and the realized scarcity value on-peak, thereby allocating supply efficiently in real time (Vickrey, 1971; Schweppe et al., 1988).

[7]Once one moves away from the efficiency benchmark toward actual tariff design and policy adoption, distributional conflict quickly becomes central. In electricity, Joskow and Wolfram (2012) stresses that dynamic pricing creates winners and losers and that fear of large redistributions is a major obstacle to adoption, while Cahana et al. (2022) shows that real-time pricing can be regressive. In transport, Hall (2021) estimates the distributive trade-offs associated with alternative road-toll designs. For a broader discussion of equity concerns in congestion pricing, see Ecola and Light (2009).

[8]A related literature studies public-utility tariff design with redistributive objectives outside the core peak-load problem. Feldstein (1972a,b) show that distributional equity invalidates the standard Ramsey–Boiteux rule. This insight has been applied notably to electricity by Levinson and Silva (2022); Feger et al. (2022). These articles share the conclusion that redistribution requires prices to differ across consumers, but they take the tariff format as given rather than deriving the full optimal mechanism from primitives.

[9]Similarly, Räsänen et al. (1997) and Martimort et al. (2020) study related public-utility environments in which redistributive preferences interact with incentive compatibility and distort the optimal tariff structure.

regulation, rationing, subsidy design, and nonlinear pricing (Pai and Strack, 2023; Ahlvik et al., 2024; Dworczak et al., 2021; Kang, 2021; Kang and Watt, 2024; Mäkimattila, 2025; Cremer and Gahvari, 2002). It is particularly close to Akbarpour et al. (2024), which also studies the allocation of a scarce resource to agents characterized by a publicly observed label and a privately observed type. The main differences are that my environment features a divisible good subject to a hard capacity constraint and an endogenous revenue requirement. As a result, the redistributive trade-off is shaped not only by primitives but also by the equilibrium surplus generated by the constrained-optimal allocation.[10]

The main results can be expressed in terms of a small set of interpretable equilibrium objects, connecting to the local-reform approach in public finance (Saez, 2001). In particular, the role of utility curvature mirrors that of demand or labor-supply responsiveness in the sufficient-statistics literature (Chetty and Saez, 2010; Kleven, 2021). It governs how scarce capacity is reallocated among consumers as capacity expands, and how changes in redistributive preferences translate into reallocation of quantity across types.[11] The article also introduces two additional equilibrium objects specific to this environment: one governing how marginal revenue is sorted across types as capacity expands, and one governing how redistributive preference perturbations reallocate quantities across types.

Because consumers belong to observable categories while their types are privately known, the article also provides foundations for group-based differentiation, in the tradition of (Akerlof, 1978; Cremer and Gahvari, 2002). This logic is relevant for essential goods subject to capacity constraints. In electricity markets, (Borenstein, 2012) evaluates the California CARE program and studies how tariff design across household groups affects efficiency and redistribution, while Burger et al. (2020) provides related simulation work on alternative tariff designs protecting low-income consumers.[12] Where prices are not an ethically acceptable rationing device, the same logic operates through non-price instruments. For instance, reserve systems in vaccine allocation explicitly allocate scarce capacity across observable groups such as essential workers, age cohorts, or medically vulnerable populations (Akbarpour et al., 2023; Pathak et al., 2024).

[10] Kang and Watt (2024) studies a related redistributive mechanism design problem in which the correlation between demand and redistributive preferences also shapes the optimal allocation. My setting differs by featuring a hard capacity constraint and an endogenous revenue margin.

[11] In that sense, it is related in spirit to reduced-form sorting effects such as the consumption-sorting channel in recent work on carbon taxation (Bierbrauer, 2024; Ahlvik et al., 2024).

[12] Similar issues arise in water and transport, where policy debates focus on group-based protections and access rules tied to observable characteristics; see (Hjort, X.; Leflaive, 2020; Manville et al., 2022).

2 Environment and mechanism

2.1 Environment

Consumers. There is a unit mass of consumers. Each consumer is characterized by a pair (i, θ) and is subject to a common shock s unrelated to type. The index $i \in N = \{1, \dots, n\}$ denotes the consumer category (e.g., households or industrial users). Categories are publicly observed, and category i has mass $\mu_i > 0$ with $\sum_i \mu_i = 1$. Conditional on category i, each consumer has a privately observed type $\theta \in \Theta_i = [\underline{\theta}_i, \overline{\theta}_i]$, drawn from a distribution G_i with density $g_i > 0$ on Θ_i. Types are independent across consumers and are known to the consumer but not to the market designer. All consumers are subject to a common demand shock $s \in S = [\underline{s}, \overline{s}]$, drawn from a distribution F with density $f > 0$. The realization of s is observed by all agents.

Preferences. A consumer of type θ derives value from consuming quantity q in state s given by

$$\theta\, U(q,s) = \theta \int_0^q u(\tilde{q}, s)\, d\tilde{q},$$

with $U(0,s) = 0$. The marginal willingness to pay $u(q,s)$ satisfies: $u_s > 0$, $u_q < 0$. Thus, marginal utility is higher in high-demand states, and it decreases with quantity. Transfers are separable, and there are no income effects. A consumer receiving allocation $(q_i(\theta,s), t_i(\theta,s))$ obtains expected surplus

$$\mathbb{E}_s[\theta U(q_i(\theta,s),s) - t_i(\theta,s)]\,.$$

Participation requires that this surplus is nonnegative for all types:

$$\mathbb{E}_s[\theta U(q_i(\theta,s),s) - t_i(\theta,s)] \geq 0, \qquad \forall \theta \in \Theta_i, \tag{IR}$$

where $\mathbb{E}_s[\cdot]$ denotes expectation with respect to the distribution of states s.

Technology and feasibility. The market designer chooses the allocation of the good subject to capacity constraints. Let $q_i(\theta,s)$ denote the quantity allocated to a consumer of type θ in category i at state s. The total quantity allocated in state s is

$$Q(s) = \sum_i \mu_i \mathbb{E}_i[q_i(\theta,s)],$$

where $\mathbb{E}_i[\cdot]$ denotes expectation with respect to the distribution of types θ. Capacity is determined by an investment level $k \geq 0$, with cost $I(k)$, where I is increasing and convex. Production has zero marginal cost. Feasibility requires that total allocation does not exceed capacity in any state:

$$Q(s) \leq k \qquad \forall s \in S. \tag{K}$$

The designer must also satisfy a revenue requirement:

$$\sum_i \mu_i \mathbb{E}_{(s,i)}[t_i(\theta,s)] - I(k) \geq 0, \tag{R$^\times$}$$

where $\mathbb{E}_{(s,i)}[\cdot]$ denotes expectation with respect to both the distribution of types θ and the distribution of states s.

Designer's objective. The market designer maximizes a weighted expected sum of consumer surpluses. Let $\lambda_i(\theta) \geq 0$ denote the welfare weight assigned to a consumer of type θ in category i, and let $\tilde{\lambda}_i := \mathbb{E}_i[\lambda_i(\theta)]$ denote the average weight in category i. The objective is

$$\sum_i \mu_i \, \mathbb{E}_{(s,i)}[\lambda_i(\theta) \, (\theta U(q_i(\theta,s),s) - t_i(\theta,s))] \, .$$

Timing. The designer commits ex ante to a capacity level k and to a direct mechanism. Consumers privately observe their types and report them to the designer. The state s is then realized, and the mechanism determines allocations and transfers as a function of reported types and the realized state. Equivalently, the designer chooses k together with the incentive-compatible allocation and transfer rules so as to maximize ex ante welfare.

2.2 Mechanism

I now consider the allocation problem under private information. The consumer's category i is publicly observed, while the type θ is privately known to the consumer.

By the revelation principle, it is without loss to focus on direct mechanisms. A direct mechanism specifies, for each category i, a quantity schedule $q_i(\theta,s)$ and a transfer schedule $t_i(\theta,s)$. For private information, the mechanism is feasible if it satisfies the previous constraints R$^\times$, K, and IR, as well as an incentive-compatibility constraint. Truthful reporting requires that, for all $\theta, \hat{\theta} \in \Theta_i$,

$$\mathbb{E}_s[\theta U(q_i(\theta,s),s) - t_i(\theta,s)] \;\geq\; \mathbb{E}_s\left[\theta U(q_i(\hat{\theta},s),s) - t_i(\hat{\theta},s)\right]. \tag{IC$^\times$}$$

Constraint IC$^\times$ imposes incentive compatibility in expectation over the common shock s. Under IC$^\times$, transfers can be chosen such that the expected surplus of a consumer of type θ in category i satisfies

$$\mathbb{E}CS_i(\theta) := \mathbb{E}_s[\theta U(q_i(\theta,s),s) - t_i(\theta,s)] = \mathbb{E}\underline{CS}_i + \int_{\underline{\theta}_i}^{\theta} \mathbb{E}_s\left[U\big(q_i(\tilde{\theta},s),s\big)\right] \, d\tilde{\theta}. \tag{$\mathbb{E}CS_i^\times$}$$

for some boundary term $\mathbb{E}\underline{CS}_i \geq 0$. This representation follows from standard envelope arguments.[13]

[13] In general, IC in expectation is weaker than pointwise IC: monotonicity of the expected marginal payoff does not, by itself, imply that the allocation is nondecreasing in θ for every state s separately. I verify ex post that the optimal solution satisfies pointwise IC.

Virtual-surplus primitives. Let $\gamma_i(\theta)$ denote the inverse hazard rate of G_i and $J_i(\theta)$ the virtual surplus function associated with a consumer from category i and type θ

$$\gamma_i(\theta) := \frac{1 - G_i(\theta)}{g_i(\theta)} \quad \text{and} \quad J_i(\theta) := \theta - \gamma_i(\theta).$$

Importantly, I do not exclude $J_i(\theta) < 0$. A negative J_i means that the surplus that can be raised from after allocating a unit of the consumer is lower than the information rent conceded to higher types. I maintain throughout the following regularity condition

Assumption Reg. *For each category i, the virtual surplus $J_i(\theta)$ is strictly increasing on $[\underline{\theta}_i, \overline{\theta}_i]$.*

Then define $\Lambda_i(\theta) := \gamma_i(\theta)\mathbb{E}_i[\lambda_i(\tilde{\theta}) \mid \tilde{\theta} \geq \theta]$. This term captures the social value of the information rents accruing to all types $\tilde{\theta} > \theta$ when type θ receives a marginal unit. Incentive compatibility forces the designer to also compensate all higher types, at a rate governed by the inverse hazard rate $\gamma_i(\theta)$, and these rents are evaluated at the average welfare weight of types above θ.

Lemma 1 (Virtual surplus representation)**.** *The program of the market designer can be expressed as:*

$$\max_{q_i(\theta,s),\, \mathbb{E}\underline{CS}_i} \quad \sum_{i \in N} \mu_i \left\{ \tilde{\lambda}_i \; \mathbb{E}\underline{CS}_i + \mathbb{E}_{(s,i)} \left[\; \Lambda_i(\theta) \; U(q_i(\theta,s),s) \; \right] \right\} \tag{CS$^\times$}$$

$$s.t. \quad \sum_{i \in N} \mu_i (\mathbb{E}_{(s,i)} \left[U(q_i(\theta,s),s) J_i(\theta) \right] - \mathbb{E}\underline{CS}_i) \; - \; I(k) \; \geq \; 0. \tag{R}$$

with the same capacity constraint K given allocation schedule $q_i(\theta, s)$ and the participation constraint for the lower type ($\mathbb{E}\underline{CS}_i \geq 0$).

Proof. See Appendix C.1 □

The lemma separates the transfer and allocation problems. When one of the budget or IR constraints binds while the other is slack, a lump-sum transfer can be used to relax the binding constraint without violating the other.

3 The Constrained Peak-Load Mechanism

This section characterizes the constrained mechanism for a single category and introduces the objects used in the rest of the analysis. The analysis proceeds in two parts. The first asks how a capacity expansion reshapes surplus across types. The second holds capacity fixed and asks how a change in redistributive preferences reallocates quantities and surplus.

Off-peak, the allocation is independent of θ. On-peak, Assumption Int ensures that the effective social weight is strictly increasing in θ, which implies that the allocation is strictly increasing in θ in every binding state simultaneously. The optimal allocation, therefore, satisfies pointwise IC, and imposing IC in expectation is without loss at the optimum.

3.1 Characterization

I begin with the basic structure of the optimal allocation for a single category. Following Lemma 1, the problem selects an allocation rule $q_i(\theta, s)$ to maximize category i's contribution to the designer's objective, subject to a category-specific capacity level k_i and revenue requirement I_i. Since transfers enter consumer surplus with a negative sign, any slack in the revenue requirement can be rebated in a lump sum to consumers while preserving incentive compatibility, raising the objective. Therefore the revenue requirement binds at the optimum: $\mathbb{E}_{(s,i)}[t_i(\theta, s)] = I_i$. Substituting the envelope representation $\mathbb{E}CS_i^{\times}$ into this equality and integrating by parts with respect to θ, using the definition of $\gamma_i(\theta)$ yields the single combined constraint

$$\mathbb{E}\underline{CS}_i = \mathbb{E}_{(s,i)}[U(q_i(\theta, s), s)\, J_i(\theta)] - I_i. \qquad (\mathbb{E}\underline{CS}_i)$$

This reformulation combines the participation and budget requirements in a single constraint when the designer maximizes consumer surplus. Throughout the article, I refer to this expression as the available surplus from the mechanism, which is redistributed via a uniform lump sum transfer to the consumers. Individual surpluses are then

$$\mathbb{E}CS_i(\theta) = \underbrace{\mathbb{E}\underline{CS}_i}_{\text{fixed transfer}} + \underbrace{\int_{\underline{\theta}_i}^{\theta} \mathbb{E}_s\left[U\left(q_i(\tilde{\theta}, s), s\right)\right] d\tilde{\theta}}_{\text{IC transfer}}. \qquad (\mathrm{ECS}_i)$$

Substituting $\mathbb{E}\underline{CS}_i$ into the objective, the problem becomes

$$\max_{q_i(\theta,s)} \quad \mathbb{E}_{(s,i)}\left[\Gamma_i(\theta)\, U(q_i(\theta, s), s) - \tilde{\lambda}_i I_i\right] \qquad (\mathrm{CS})$$

$$\text{s.t.} \quad \mathbb{E}_{(s,i)}[U(q_i(\theta, s), s)\, J_i(\theta)] - I_i \geq 0, \qquad (\text{IR-R}_i)$$

$$\mathbb{E}_i[q_i(\theta, s)] \leq k_i \qquad \forall s. \qquad (\mathrm{K}_i)$$

where $\Gamma_i(\theta) := J_i(\theta)\, \tilde{\lambda}_i + \Lambda_i(\theta)$. Let $\varepsilon_i(s) \geq 0$ and $\beta_i \geq 0$ denote the multipliers on the capacity constraint K_i and the IR/budget constraint IR-R$_i$.

The first-order condition is

$$u(q_i(\theta, s), s)\ \mathcal{G}_i(\theta) - \varepsilon_i(s) = 0, \qquad (\mathrm{FOC}_q)$$

where $\mathcal{G}_i(\theta) := \Gamma_i(\theta) + J_i(\theta)\, \beta_i$ is the effective weight of a type θ. It captures the total social return from allocating a marginal unit to type i: the revenue channel $\tilde{\lambda} J_i$ values the available surplus J_i that can be redistributed to all consumer, so it is valued at the average welfare weight, Λ_i captures the redistributive value of the rent conceded to higher types, and βJ_i reflects the tightness of the budget constraint. When the budget is slack, $\beta = 0$ and $\mathcal{G}_i$ reduces to Γ_i. I focus on interior solutions and impose two regularity conditions. The first discipline of the allocation rule: it ensures all types receive a strictly positive allocation and that the monotonicity constraint is automatically satisfied at the optimum.[14]

Assumption Int. *For all $\theta \in [\underline{\theta}_i, \overline{\theta}_i]$, $\mathcal{G}_i(\theta) > 0$ and $\mathcal{G}_i(\theta)$ is non-decreasing in θ, so that the monotonicity constraint on $q_i(\theta, s)$ is satisfied at the optimum. Moreover, the marginal utility function satisfies $\lim_{q\to 0} u(q,s) = +\infty$.*

Define $S_i = [\underline{s}, s_i)$ and $T_i = [s_i, \overline{s}]$ as the off-peak and on-peak sets under the optimal allocation, where s_i satisfies $\mathbb{E}_i[q_i(\theta, s_i)] = k_i$.

Lemma 2 (Off-peak and on-peak allocation)**.** *Assume Int holds and that the problem is convex. Under the weighted objective with type-dependent weights $\lambda_i(\theta)$, the allocation is characterized by a unique quantity schedule $q_i(\theta, s)$ for each θ and s:*

$$u(q_i(\theta,s),s) = 0 \quad s \in S_i, \qquad \text{and} \qquad u(q_i(\theta,s),s)\ \mathcal{G}_i(\theta) = \varepsilon_i(s) \quad s \in T_i. \tag{SB-S}$$

Proof. See Appendix C.2 □

3.2 Sufficient statistics and assumptions

I now define the sufficient statistics that drive the main results of the article. I start with the capacity-sorting margin and then discuss the preference-sorting margin.

3.2.1 Capacity margin

Define

$$R_q(\theta) := -\frac{u(q_i(\theta,s),s)}{u_q(q_i(\theta,s),s)},$$

The object R_q is determined by the curvature of demand and is the reciprocal of the coefficient of absolute risk aversion, with $R_q(\theta) \geq 0$ since $u_q < 0$. It governs how an increase in k is distributed across types at the optimal allocation. To recover the object, differentiating FOC_q at a fixed budget multiplier β_i and equating across two types θ_a, θ_b yields $dq_i(\theta_a)/R_q(\theta_a) = dq_i(\theta_b)/R_q(\theta_b)$.

Next, define

$$R_\theta(\theta) := \frac{J_i(\theta)}{\mathcal{G}_i(\theta)},$$

as the ratio of a type's ability to raise the available surplus (J_i) to its effective social weight ($\mathcal{G}_i$), and measures the marginal revenue generated per unit of social value when serving type i. In particular, it governs how

[14] Appendix A.2.1 shows that with three types, monotonicity of effective weights holds if and only if redistributive preferences are not too concentrated on any single type. If monotonicity fails, standard Myerson-style ironing restores it by pooling types in the non-monotone region. Over any ironed interval, the optimal allocation is constant across types, so the sorting margins in the article vanish rather than reverse, and their effect depends on the non-ironed types. This stands in contrast to settings where the supply side independently restricts optimal mechanisms, either through indivisibility, as in Condorelli (2013), or through a fixed quality schedule, as in Akbarpour et al. (2024).

the marginal virtual surplus $u(q_i,s)J_i(\theta)$ varies across types. If $R_\theta(\theta)$ is increasing in θ, then the marginal virtual surplus tends to be lower for lower types.

Define the weights

$$w_i(\theta,s) := -\frac{1}{u_q(q_i(\theta,s),s)\,\mathcal{G}_i(\theta)}, \quad \overline{w}_i(s) := \mathbb{E}_i[w_i(\theta,s)]\,.$$

These weights measure the responsiveness of q_i to the scarcity rent. To recover the object, differentiating FOC_q gives the equality $R_q(\theta) = w_i(\theta,s)\varepsilon_i(s)$, so w_i governs how the optimal allocation responds to a change in capacity.

Now, combine the previous objects to define

$$\mathbb{E}_i^w[uJ \mid s] := \frac{1}{\overline{w}_i(s)}\mathbb{E}_i[w_i(\theta,s)\,u(q_i(\theta,s),s)\,J_i(\theta)] = \frac{1}{\overline{w}_i(s)}\mathbb{E}_i\left[R_\theta(\theta)R_q(\theta)\right],$$

where $\mathbb{E}_i^w[\cdot \mid s]$ denote the average weighted by $w_i(\theta,s)$. Average over binding states yields

$$\mathbb{E}_{(s,i)}^w[uJ \mid s] := \mathbb{E}_s\left[\mathbb{E}_i^w[uJ \mid s]\mathbf{1}_{\{s\in T_i\}}\right], \tag{EuJ}$$

Where $1_{\{s\in T_i\}}$ is the indicator of the on-peak period. It depends on k_i, but as the boundaries effect are null, I drop the reference to k_i. The sign of this weighted value is the *capacity-sorting margin*. It admits two economic interpretations that operate under different regimes. When the budget constraint is slack ($\beta_i = 0$), it has the same sign as the derivative of the available surplus $\mathbb{E}\underline{CS_i}$ with respect to k_i. When the budget constraint binds ($\beta_i > 0$), its sign determines whether a capacity expansion loosens or tightens the budget, via $\frac{\partial \beta_i}{\partial k_i}$. Both interpretations are established formally in Theorem 1 and illustrated in the two-type case of Section 3.3. The two expressions of $\mathbb{E}_i^w[uJ \mid s]$ above give two interpretations of the capacity-sorting margin. The first interpretation, reading the middle expression, is that $R_q(\theta_i)$ governs how a capacity expansion is distributed across types, while $R_\theta(\theta_i)$ governs the marginal revenue generated by each type's allocation. When R_q is decreasing in θ, a capacity expansion flows disproportionately to low types; when $R_\theta(\theta)$ is also increasing in θ, those are precisely the types that contribute least to available surplus, so expansion reduces this surplus and tightens the budget constraint. The second interpretation, reading the right-hand expression, is that $\mathbb{E}_i^w[uJ \mid s]$ is a w_i-weighted average of marginal virtual surplus $u(q_i)J_i$, where the weights w_i jointly reflect the responsiveness of each type's demand and its effective social value.

To obtain characterization, I impose the following regularity condition.

Assumption Monot.. *$R_q(\theta)$ and $R_\theta(\theta)$ are both monotone on $[\underline{\theta}_i, \overline{\theta}_i]$. Moreover, $R_q(\theta)$ exhibits log-monotone differences in (θ, k_i).*

In particular, I show in Lemma 6 of the Appendix that EuJ is single-crossing in k_i when Assumption Monot.. Hence, the available surplus to be redistributed is either single-peaked or single-dipped.[15] I discuss in section A.1 in the Appendix, the conditions for establishing the monotonicity of those ratios.

Finally, a useful specification used in the article is CARA utility as a consumer's responsiveness is identical regardless of the quantity allocated.

Assumption CARA. *$U(q,s) = v(s)/\alpha(1-\exp(-\alpha q))$ for some $\alpha > 0$ and some function $v(s) > 0, v_s(s) > 0$.*

Under Assumption CARA, $R_q = 1/\alpha$ is constant: a consumer's responsiveness is identical regardless of the quantity allocated. Additional capacity is therefore distributed across types in proportions determined solely by welfare weights, and the behavioral channel through which a preference shift distorts the allocation across types vanishes.

3.2.2 Preference margin

I start by defining a perturbation of redistributive preferences. Let $\Delta\lambda_i(\theta)$ denote the marginal change in redistributive preferences.[16] The induced change in the effective social weight is

$$\Delta\mathcal{G}_i(\theta) := J_i(\theta)\big(\Delta\tilde{\lambda}_i + \Delta\beta_i\big) + \Delta\Lambda_i(\theta), \quad \Delta\Lambda_i(\theta) := \gamma_i(\theta)\mathbb{E}_i\big[\Delta\lambda_i(\tilde{\theta}) \mid \tilde{\theta} \geq \theta\big], \quad \Delta\tilde{\lambda}_i = \mathbb{E}_i[\Delta\lambda_i(\theta)],$$

where $\Delta\beta_i$ is the induced change in the budget multiplier at the new optimum. The key statistic is the proportional change in the effective social weight of type θ,

$$R_{\mathcal{G},i}(\theta) := \frac{\Delta\mathcal{G}_i(\theta)}{\mathcal{G}_i(\theta)},$$

which measures how the perturbation tilts the effective weight of type θ relative to the baseline allocation. This is the preference-sorting margin, and it governs how a change in redistributive preferences reshapes the constrained allocation at fixed capacity. This object appears directly when differentiating FOC$_q$. Substituting $u(q_i(\theta,s),s) = \varepsilon_i(s)/\mathcal{G}_i(\theta)$ into the differentiated first-order condition yields

$$\frac{\Delta q_i(\theta,s)}{R_q(\theta)} = R_{\mathcal{G},i}(\theta) - \frac{\Delta\varepsilon_i(s)}{\varepsilon_i(s)}. \tag{ΔFOC$_q$}$$

[15] The economic interpretation of log-monotone differences in (θ, k_i) is intuitive. Suppose that, for each k_i, $R_\theta(\theta)$ is increasing in θ, so higher types receive higher effective virtual surplus, and that $R_q(\theta)$ is log-submodular in (θ, k_i). Then capacity expansion disproportionately raises R_q for lower types, in the sense that for $k_i' > k_i$, the ratio $R_q(\theta')/R_q(\theta)$ is decreasing in θ. Thus, as capacity expands, the weighted average EuJ shifts monotonically from the regime in which the average effect dominates toward the regime in which the sorting effect dominates.

[16] I remain agnostic about whether a perturbation $\Delta\lambda_i(\theta)$ represents a shift in the designer's social welfare function or a change in the joint distribution of the latent welfare characteristic λ and the preference type θ conditional on category i. The former corresponds to a change in the designer's normative priorities; the latter to updated beliefs about the correlation between willingness to pay and social need. The comparative statics apply under either interpretation.

A preference perturbation affects the optimal allocation through two forces. The first is the proportional adjustment of the scarcity rent, $\Delta\varepsilon_i(s)/\varepsilon_i(s)$, which is common to all types within state s and captures how tight capacity becomes once preferences have shifted. The second is the proportional revaluation of each type's effective social weight, $R_{\mathcal{G},i}(\theta)$, which captures how the perturbation changes the social value of serving type θ relative to the baseline allocation. At the category-state level, the relevant aggregate object is the weighted average $\mathbb{E}_i^w[R_{\mathcal{G}} \mid s]$, which summarizes the average preference tilt across types using the weights $w_i(\theta, s)$ defined above.

Appendix Lemma 7 provides primitive conditions on preference perturbations under which $R_{\mathcal{G}}(\theta)$ has a constant sign across types. The first part of the Lemma states that a proportional tilt of welfare weights toward higher types shifts the social value of the information rent given to types above θ (i.e., $\Lambda_i(\theta)$) upward relative to the baseline across the type space.[17] The second part connects the direction of a preference shift to the sign of $R_{\mathcal{G}}(\theta)$, and requires two conditions. The effective weight must not decrease on average, and the average preference shift must be at least as large as the effective-weight shift after accounting for the budget shadow cost. These conditions simplify under two special cases on the equilibrium object β_i, which will be used later for the across-category results. Tilting welfare weights proportionally toward higher types, while not decreasing the average weight, yields $R_{\mathcal{G}}(\theta) > 0$ for all θ; and conversely, shifting proportionally toward lower types while not increasing the average weight yields $R_{\mathcal{G}}(\theta) < 0$ for all θ. A standard perturbation in the redistribution literature is a (monotonic) mean-preserving rotation of $\lambda_i(\theta)$, i.e. $\Delta\tilde{\lambda}_i = 0$ (see Kang and Watt (2024) for a recent use in the context of redistributive mechanism design). Under such a rotation, $R_{\mathcal{G}}(\theta)$ has a constant sign throughout the type space. It is positive when the rotation favors higher types, and negative when it favors lower types.

3.3 A two-type illustration of the sorting margins

To illustrate the two sorting margins in the simplest possible environment, consider a single category with two types, $\theta \in \{\theta_L, \theta_H\}$, occurring with probabilities p_L and p_H, where $p_L + p_H = 1$. I focus on the peak allocation when the participation constraint is slack, so $\beta = 0$ and $\mathcal{G}_l = \Gamma_l$, suppress the state index s, and assume that $l \in \{L, H\}$. A type-θ consumer then obtains utility $\theta U(q)$ from quantity q. In particular, the example shows how the capacity-sorting margin governs the surplus incidence of a marginal capacity expansion, and how the preference-sorting margin, through the proportional change in effective social weight $R_{\mathcal{G}}(\theta)$, governs the direction of reallocation induced by a change in redistributive preferences.

[17] The first result of the Lemma has a direct implication for the sign of $R'_\theta(\theta)$. From expression signR_θ, only the IC term depends on preferences; the revenue term is determined solely by the type distribution. As a maintained example, assume $\Lambda_i^{1\prime}(\theta) < 0$, $\Lambda_i^{2\prime}(\theta) < 0$, and $J_i'(\theta) > 0$. In that case, a preference shift toward higher types favors $R'_\theta(\theta) > 0$. For types generating revenue ($J_i(\theta) > 0$), $R'_\theta(\theta) > 0$ holds under both weight functions. For types that do not generate revenue ($J_i(\theta) < 0$), the condition in signR_θ is easier to satisfy under a preference shift toward higher types, since $\Lambda_i^{\Delta\prime}/\Lambda_i^{\Delta} \geq \Lambda_i'/\Lambda_i$ by Lemma 7.

In the two-type case, the expressions for available surplus and individual surplus simplify as follows.

$$\mathbb{E}\underline{CS} := \underbrace{\sum_l p_l \theta_l U(q_l)}_{\text{surplus}} - \underbrace{\sum_l p_l \gamma_l U(q_l)}_{\text{IC rent}} - \underbrace{I}_{\text{cost}} = p_L J_L U(q_L) + p_H J_H U(q_H) - I, \tag{$\mathbb{E}\underline{CS}_{LH}$}$$

$$\mathbb{E}CS(\theta_L) = \underbrace{\mathbb{E}\underline{CS}}_{\text{fixed transfer}}, \qquad \mathbb{E}CS(\theta_H) = \mathbb{E}\underline{CS} + \underbrace{(\theta_H - \theta_L)U(q_L)}_{\text{IC transfer}}. \tag{$\mathbb{E}\underline{CS}_{LH}$}$$

Hence the low type receives the full redistributable surplus as a lump-sum transfer, while the high type receives the same lump-sum transfer plus a type-dependent informational rent.

The inverse hazard rate of the low type is $\gamma_L = p_H/p_L$, while $\gamma_H = 0$ since no information rent accrues above the high type. The virtual terms are

$$J_L = \theta_L - \frac{p_H}{p_L}(\theta_H - \theta_L), \qquad J_H = \theta_H.$$

$J_L < 0$ arises when the low type's value lies below the high type's value weighted by the share of high types, $\theta_L < p_H \theta_H$. The social value of this information rent flowing upward to θ_H when θ_L receives a marginal unit is

$$\Lambda_L = \frac{p_H}{p_L}(\theta_H - \theta_L)\lambda_H, \qquad \Lambda_H = 0$$

Capacity sorting margin with two types. I now study how a capacity expansion affects the available surplus, $\mathbb{E}\underline{CS}_{LH}$. The two-type surplus decomposition in $\mathbb{E}\underline{CS}_{LH}$ implies that the sign of $\partial_k \mathbb{E}\underline{CS}_{LH}$ determines the individual welfare change. In particular, whenever $\partial_k \mathbb{E}\underline{CS}_{LH} < 0$, the low type loses from an increase in k regardless of the welfare weight λ_L.

Differentiating $\mathbb{E}CS(\theta_L)$ with respect to k, factorizing by $\frac{\partial q_L}{\partial k}$, ε and $1/R_q(\theta_L)$ yields

$$\frac{\partial \mathbb{E}CS(\theta_L)}{\partial k} = \frac{\partial q_L}{\partial k} \frac{\varepsilon}{R_q(\theta_L)} \Big(\underbrace{\theta_L \, p_L \frac{R_q(\theta_L)}{\Gamma_L} + \theta_H \, p_H \frac{R_q(\theta_H)}{\Gamma_H}}_{\text{available surplus}} - \underbrace{(\theta_H - \theta_L) \, p_H \frac{R_q(\theta_L)}{\Gamma_L}}_{\text{IC cost}} \Big). \tag{$\partial_k \mathbb{E}\underline{CS}_{LH}$}$$

Equivalently, the same object can be rewritten in terms of the capacity-sorting margin of Section 3.2:

$$\frac{\partial \mathbb{E}\underline{CS}}{\partial k} = \frac{1}{\overline{w}} \left(p_L \, R_q(\theta_L) \frac{J_L}{\Gamma_L} + p_H \, R_q(\theta_H) \frac{J_H}{\Gamma_H} \right) = \frac{1}{\overline{w}} \sum_l p_l \, R_\theta(\theta_l) \, R_q(\theta_l) = \frac{1}{\overline{w}} \sum_l p_l \, w_l \, u(q_l) \, J_l.$$

Increasing k has two opposite effects. It raises the available quantity and therefore the available surplus terms. However, expanding k also forces the designer to leave a larger informational rent to θ_H, which mechanically reduces the transfer t_H and therefore the available surplus $\mathbb{E}\underline{CS}$.

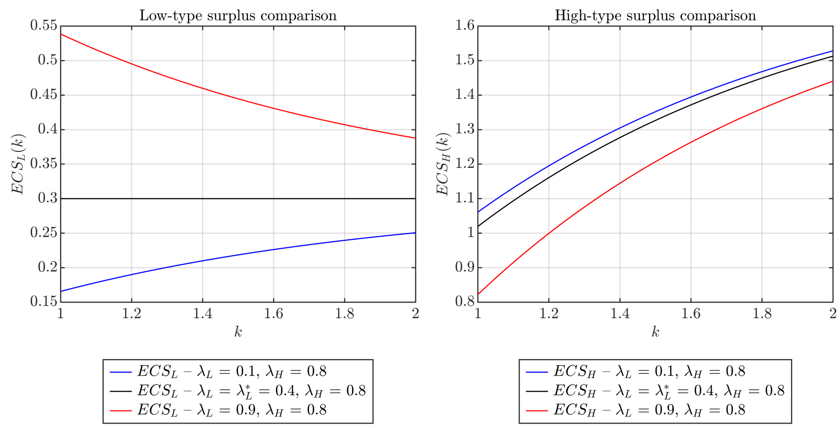

Figure 1: Individual surplus as a function of capacity, for three welfare weights $\lambda_L \in \{0.1, 0.4, 0.9\}$ with $\lambda_H = 0.8$. Left panel: low type. Right panel: high type.

A CARA specification allows to characterize clearly when the low type loses from a capacity expansion, define λ_L^* as the cutoff that solves $\partial_k \mathbb{E}\underline{CS}_{LH} = 0$. When $J_L < 0$, this cutoff is unique, and

$$\frac{\partial \mathbb{E}CS(\theta_L)}{\partial k} < 0 \qquad \Longleftrightarrow \qquad \lambda_L > \lambda_L^* = \frac{\theta_L p_H \lambda_H}{p_H(\theta_H - \theta_L) - \theta_L p_L} = -\frac{\theta_L p_H \lambda_H}{p_L J_L}.$$

Raising λ_L moves Γ_L, and the effect of Γ_L on the available surplus goes in the same direction: both channels are scaled by the low type's net contribution $p_L J_L$. The two effects reinforce each other so the net effect of λ_L on $\partial_k \mathbb{E}\underline{CS}_{LH}$ is unambiguously negative. Note that a larger θ_L raises the cutoff, making it harder to obtain $\partial_k \mathbb{E}\underline{CS}_{LH} < 0$. Intuitively, raising θ_L has two effects: a positive direct effect, through the gain in surplus from serving the low type, and the reduction in $\theta_H - \theta_L$, and a negative effect since Γ_L rises with θ_L. In the two-type case, the positive effect always dominates. Figure 1 illustrates individual surplus for both types as a function of capacity k, under three values of the welfare weight λ_L: below, at, and above the cutoff λ_L^*.

Preference-sorting margin with two types. I now show that a change in redistributive preferences need not reallocate quantity toward the type whose welfare weight rises. Consider a perturbation of redistributive preferences $(\Delta\lambda_L, \Delta\lambda_H)$, inducing a change in the average weight $\Delta\tilde{\lambda} = p_L\,\Delta\lambda_L + p_H\,\Delta\lambda_H$. The induced change in effective social weights is

$$\Delta\Gamma_L = J_L\,\Delta\tilde{\lambda} + \frac{p_H}{p_L}\,(\theta_H - \theta_L)\,\Delta\lambda_H, \qquad \Delta\Gamma_H = \theta_H\,\Delta\tilde{\lambda}.$$

These expressions make the two forces transparent. The first term, $\Delta\tilde{\lambda} J_l$, is the revenue motive: it revalues each type's contribution to available surplus at the new average weight. The second term is the redistributive motive: it revalues the informational rent flowing upward to higher types. Since no rent accrues above the high type, $\Delta\Lambda_H = 0$, so $\Delta\Gamma_H$ depends only on the average perturbation. By contrast, $\Delta\Gamma_L$ contains both terms, because serving the low type also changes the informational rent left to θ_H.

Two special cases follow directly. If only λ_L changes, then $\Delta\tilde{\lambda} = p_L\,\Delta\lambda_L$ and $\Delta\lambda_H = 0$, so $\Delta\Gamma_H = p_L\theta_H\,\Delta\lambda_L$ and $\Delta\Gamma_L = p_L J_L\,\Delta\lambda_L$. The redistributive motive is then absent, and the perturbation operates purely through the revenue channel for both types. Under a mean-preserving perturbation, instead, $\Delta\tilde{\lambda} = 0$ and $\Delta\lambda_H = -(p_L/p_H)\,\Delta\lambda_L$, so $\Delta\Gamma_H = 0$ and $\Delta\Gamma_L = -(\theta_H - \theta_L)\,\Delta\lambda_L$. In that case, the high type's effective weight is unchanged and the entire perturbation operates through the redistributive motive of the low type alone.

Using $\Delta\mathrm{FOC}_q$ for L and H and subtracting gives

$$\Gamma_L u_q(q_L)\,\Delta q_L - \Gamma_H u_q(q_H)\,\Delta q_H = \varepsilon\left(R_{\mathcal{G}}(\theta_H) - R_{\mathcal{G}}(\theta_L)\right). \qquad (\Delta q_{LH})$$

Differentiating the capacity constraint yields $p_L\,\Delta q_L + p_H\,\Delta q_H = 0$, so that $\Delta q_H = -(p_L/p_H)\,\Delta q_L$ and the two quantities always move in opposite directions. Hence the allocation shifts toward the type whose effective weight experiences the larger proportional increase:

$$\operatorname{sign}(\Delta q_L) = \operatorname{sign}\left(R_{\mathcal{G}}(\theta_L) - R_{\mathcal{G}}(\theta_H)\right), \qquad \operatorname{sign}(\Delta q_H) = \operatorname{sign}\left(R_{\mathcal{G}}(\theta_H) - R_{\mathcal{G}}(\theta_L)\right).$$

The two special cases then yield the following implications. When $\Delta\lambda_H = 0$,

$$R_{\mathcal{G}}(\theta_L) - R_{\mathcal{G}}(\theta_H) = \frac{J_L\,p_L\,\Delta\lambda_L}{\Gamma_L} - \frac{p_L\,\Delta\lambda_L}{\tilde{\lambda}} = -\frac{p_H(\theta_H - \theta_L)\lambda_H}{\Gamma_L\tilde{\lambda}}\,\Delta\lambda_L,$$

so an increase in λ_L alone lowers q_L. By contrast, when $\Delta\tilde{\lambda} = 0$,

$$R_{\mathcal{G}}(\theta_L) - R_{\mathcal{G}}(\theta_H) = -\frac{\theta_H - \theta_L}{\Gamma_L}\,\Delta\lambda_L,$$

so a mean-preserving tilt toward the high type lowers q_H and raises q_L. When only λ_L rises, the perturbation works through the revenue motive, which raises the high type's effective weight proportionally more because it enters through θ_H rather than the smaller object J_L. Under a mean-preserving tilt, instead, the revenue motive is absent, so only the redistributive motive remains, which raises Γ_L while leaving Γ_H unchanged.

More generally, combining the expressions for $\Delta\Gamma_L$ and $\Delta\Gamma_H$ and simplifying yields

$$R_{\mathcal{G}}(\theta_L) - R_{\mathcal{G}}(\theta_H) = \frac{p_H(\theta_H - \theta_L)}{\tilde{\lambda}\,\Gamma_L}\left(\lambda_L\,\Delta\lambda_H - \lambda_H\,\Delta\lambda_L\right).$$

Therefore,

$$\operatorname{sign}(\Delta q_L) = \operatorname{sign}\left(\Delta\left(\frac{\lambda_H}{\lambda_L}\right)\right), \qquad \operatorname{sign}(\Delta q_H) = -\operatorname{sign}\left(\Delta\left(\frac{\lambda_H}{\lambda_L}\right)\right).$$

Thus, the allocation shift depends on the change in the ratio λ_H/λ_L: the low type receives more quantity if and only if the relative increase in welfare weight is stronger for the high type. This ratio indicates which of the two forces dominates: the redistributive motive, which tilts the allocation toward the low type, or the revenue motive, which tilts it toward the high type. The Lemma 7 in the Appendix shows that this ratio characterization extends to the continuous-type case.

The two-type example makes the two sorting margins transparent, but some of its properties are specific to the binary environment. In particular, both the effective weight $\mathcal{G}_l$ and the ratio $R_\theta(\theta_l)$ are always increasing in θ with two types, so Assumptions Int and Monot. hold automatically and play no role. Appendices A.2.1 and A.2.2 show that both monotonicities can fail with three types. I now turn to the general analysis, which does not rely on these special features.

3.4 Redistributive incidence of capacity expansion

I now formally study the following question: who benefits from a marginal increase in capacity? Even though a larger k_i relaxes scarcity, the surplus gains from this relaxation need not accrue to all types once incentive constraints and revenue extraction are taken into account. This subsection shows that the incidence of capacity expansion is pinned down by the sign of EuJ, the capacity-sorting margin, together with the ordering of types induced by $R_\theta(\theta)$.

The individual surplus change follows from $\mathbb{E}CS_i^\times$

$$\frac{\partial \mathbb{E}CS_i(\theta)}{\partial k_i} = \underbrace{\frac{\partial \mathbb{E}\underline{CS}_i}{\partial k_i}}_{\text{fixed transfer}} + \underbrace{\int_{\underline{\theta}_i}^{\theta} \mathbb{E}_s\left[\frac{\partial q_i(\tilde{\theta},s)}{\partial k_i} u\big(q_i(\tilde{\theta},s),s\big)\right] d\tilde{\theta}}_{\text{IC transfer}}. \tag{$\partial_k CS$}$$

The derivative of the type-independent component is obtained from $\mathbb{E}CS_i^\times$ while differentiating the first-order condition FOC_q with respect to k_i gives

$$\frac{\partial q_i(\theta,s)}{\partial k_i} = w_i(\theta,s)\left(\frac{1}{\overline{w}_i(s)} - \frac{\partial \beta_i}{\partial k_i}\Big(\mathbb{E}_i^w[uJ \mid s] - u(q_i(\theta,s),s)\, J_i(\theta)\Big)\right). \tag{$\partial_k q_i$}$$

Two observations: when IR-R$_i$ binds (the IR-constrained allocation), the boundary term is null: the entire surplus change is then carried by the IC rent in $\partial_k CS$. On the other hand, when IR-R$_i$ is slack (the IR-unconstrained allocation), then the second term in $\partial_k q_i$ is null as $\beta_i = 0$. I assume the following so that the single-crossing structure is preserved after aggregation over states.

Assumption Multi. *$u(q,s) = \kappa(q)\zeta(s)$ for some functions κ and ζ with $\kappa_q < 0$ and $\zeta_s > 0$.*

Under Assumption Multi, $R_q = -\kappa(q)/\kappa_q(q)$ depends only on q and not on s: a consumer's responsiveness is identical regardless of the realization of the common shock.

Theorem 1 (Redistributive effects of capacity expansion)**.** *Fix category i, a budget requirement $I_i > 0$ and assume Monot. and Multi hold. Define the set of cutoffs*

$$\tilde{\Theta}_i := \left\{\tilde{\theta}_i \in [\underline{\theta}_i, \overline{\theta}_i] : \partial \mathbb{E}CS_i(\tilde{\theta}_i)/\partial k_i = 0\right\}.$$

IR-unconstrained allocation. *If EuJ is positive, then an increase in k_i increases expected surplus for every type θ. If EuJ is negative, $\tilde{\Theta}_i$ contains at most a unique cutoff; and if it exists expected surplus decreases for all $\theta < \tilde{\theta}_i$ and increases for all $\theta > \tilde{\theta}_i$.*

IR-constrained allocation. *The lowest type always lies on the boundary: $\underline{\theta}_i \in \tilde{\Theta}_i$. If $R_\theta(\theta)$ is increasing, a second cutoff $\tilde{\theta}_i \in (\underline{\theta}_i, \overline{\theta}_i)$ may exist and when it does:*

(i) if EuJ is positive, expected surplus increases for all $\theta \in (\underline{\theta}_i, \tilde{\theta}_i)$ and decreases for all $\theta \geq \tilde{\theta}_i$;

(ii) if EuJ is negative, expected surplus decreases for all $\theta \in (\underline{\theta}_i, \tilde{\theta}_i)$ and increases for all $\theta \geq \tilde{\theta}_i$.

If $R_\theta(\theta)$ is decreasing, the inequalities are reversed. If no interior cutoff exists, expected surplus is uniformly positive (resp. negative) across all $\theta \in (\underline{\theta}_i, \overline{\theta}_i)$.

Proof. See Appendix C.4 □

Expanding capacity is not always surplus-improving for all consumers, and the identity of who gains/loses depends on the regime and on the joint sign of EuJ and $R_\theta(\theta)$. Importantly, the mechanism through which a capacity expansion determines the progressivity/regressivity of the capacity expansion is not the same across regimes. Under the IR-unconstrained regime, it operates through changes in the available surplus $\mathbb{E}\underline{CS}_i$, while under the IR-constrained regime it operates through changes in the optimal allocation $q_i(\theta, s)$.

When the participation constraint is slack ($\beta_i = 0$), the optimal allocation is increasing in k_i for all types in binding states. Hence, the IC component is always increasing. However, the sign of the type-independent component in $\partial_k CS$ depends entirely on the sign of EuJ, which therefore summarizes the marginal redistributive effect of expanding capacity. If EuJ>0, both the common surplus component and the induced IC rents increase with k_i, so all types gain. If EuJ<0, lower types lose while higher types gain, with a unique cutoff as stated in Theorem 1.

When the participation constraint binds, the fixed transfer is null, so the surplus effect of a capacity expansion is driven by the quantity margin. However, the monotonicity of quantities in k_i may fail because a capacity expansion also changes the tightness of the financing constraint through the endogenous response of β_i, as shown in $\partial_k q_i$. The bracketed term in that expression is the deviation of a type's marginal virtual surplus $u(q_i(\theta,s),s)J_i(\theta)$ from its w_i-weighted average; its sign is therefore determined by the relative ranking of types according to marginal virtual surplus, which is captured by $R_\theta(\theta)$. The sign of $\partial\beta_i/\partial k_i$ is again governed by EuJ, which captures how capacity expansion tightens or loosens the budget constraint. Therefore, the condition on EuJ and $R_\theta(\theta)$ has a direct interpretation: suppose that expanding k_i makes the budget less tight, so that β_i decreases in k_i. The need for revenue extraction then falls, and the designer lowers the marginal weight on types that generate higher marginal virtual surplus.

Assumption Monot. allows me to characterize cleanly the set of winners and losers from a capacity expansion. Without this requirement, the model may exhibit a richer set of cutoffs. In Appendix Proposition 4, I show that EuJ also governs the transition between the IR-unconstrained and IR-constrained regimes. This transition depends on whether the available surplus is single-peaked or single-dipped. Figure 2 summarizes this regime geometry by illustrating how the distributional incidence of a capacity expansion changes across regimes and across capacity levels, while the section A.3 in the Appendix provides the full characterization. A direct implication is that the redistributive effect of additional capacity is not invariant to the installed capacity level: the same marginal expansion can be progressive at low k_i and regressive at high k_i. I also provide a closed-form numerical simulation in the Appendix Figure 9 illustrating how the identity of the winners and losers from a capacity expansion depends on the current level of capacity.

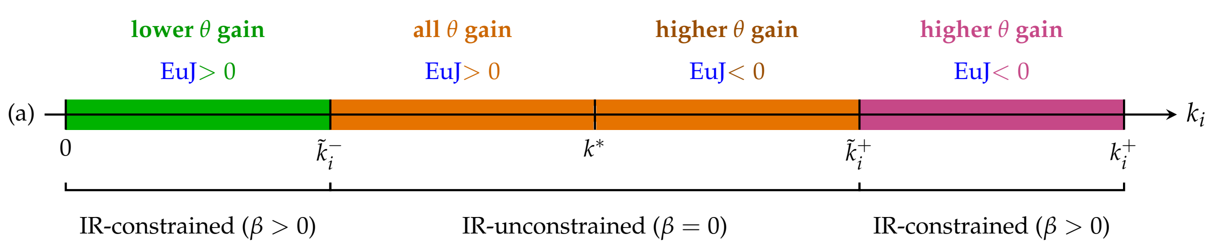

Figure 2: Joint characterization of the implementation regime and surplus effects of capacity expansion. Each color represents a specific redistributive effect of capacity expansion across different capacity levels. The example represents a single-peaked revenue (submodular R_q and increasing $R_\theta(\theta)$), so EuJ crosses from $+$ to $-$.

3.5 Preference margin and optimal allocation

The previous section asked how a given redistributive objective shapes the surplus consequences of capacity expansion. I now reverse the question: holding k_i fixed, how does a change of redistributive preferences reshape the optimal allocation? This allows me to study how incentive constraints shape the incidence of redistributive preference changes for a given capacity constraint, via the effective social weights.[18]

[18] This comparative static is also the building block of the across-category analysis in Section 4. Comparing the optimal allocations of two categories that differ marginally in preferences is equivalent to evaluating the local effect of a preference perturbation at fixed capacity.

Lemma 3 (Reallocation under preference perturbations)**.** *Fix category i and an across-category allocation (k_i, I_i). For any local perturbation $\Delta\lambda_i(\theta)$ and any allocation regime,*

$$\Delta q_i(\theta, s) = R_q(\theta)(R_{\mathcal{G}}(\theta) - \mathbb{E}_i^w[R_{\mathcal{G}}(\theta) \mid s])\,, \qquad s \in T_i.$$

Proof. See Appendix C.5 □

To interpret Lemma 3, type θ receives more quantity if and only if its proportional change in effective social weight exceeds the proportional average. The intuition is the following. From ΔFOC_q, the linearization of the pointwise first-order condition yields

$$\Delta q_i(\theta, s) = w_i(\theta, s)\big(u(q_i, s)\,\Delta\mathcal{G}_i(\theta) - \Delta\varepsilon_i(s)\big) = R_q(\theta)(R_{\mathcal{G}}(\theta) - \mathbb{E}_i^w[R_{\mathcal{G}}(\theta) \mid s])\,. \tag{Δq}$$

That is the condition of receiving more or less allocation is determined by whether the preference change raises the marginal utility valued at the effect social weights of type θ by more or less than the equilibrium adjustment in the scarcity rent.[19] Since capacity is fixed, the reallocation is zero-sum across types within each binding state: $\mathbb{E}_i[\Delta q_i(\theta, s)] = 0$ for each $s \in T_i$.

The sorting effect of the optimal allocation induced by a preference perturbation is different than the direction of the perturbation itself. Appendix Proposition 5 shows that the behavioral response $\Delta q_i(\cdot, s)$ can be non-monotone in types, so a tilt of preferences toward higher or lower types does not mechanically translate into a monotone quantity reallocation. This non-monotonicity of the optimal allocation stems from the fact that the element that drives its response is the proportional change $R_{\mathcal{G}}(\theta)$ and not the absolute change $\Delta\mathcal{G}_i(\theta)$. While a change in preferences may induce a monotonic variation in the absolute change $\Delta\mathcal{G}_i(\theta)$ along types, this need not be the case for the proportional change $R_{\mathcal{G}}(\theta)$. To see this, consider the extreme types $\underline{\theta}_i$ and $\overline{\theta}_i$ and a perturbation of an optimal allocation in the IR-unconstrained regime. At $\underline{\theta}_i$, the social value of the information rent is the information rent valued at the expected welfare weights: $\gamma_i(\theta)\tilde{\lambda}_i$, so the second and third terms cancel: $\mathcal{G}_i(\underline{\theta}_i) = \underline{\theta}_i\tilde{\lambda}_i$. At $\overline{\theta}_i$, the information rent is absent: $\gamma_i(\theta) = 0$, so the second and third terms are null: $\mathcal{G}_i(\overline{\theta}_i) = \overline{\theta}_i\tilde{\lambda}_i$. The perturbation yeilds similar values: $\Delta\mathcal{G}_i(\underline{\theta}_i) = \underline{\theta}_i\Delta\tilde{\lambda}_i$ and $\Delta\mathcal{G}_i(\overline{\theta}_i) = \overline{\theta}_i\Delta\tilde{\lambda}_i$. Therefore, while I can have $\Delta\mathcal{G}_i(\underline{\theta}_i) \neq \Delta\mathcal{G}_i(\overline{\theta}_i)$, I always have $R_{\mathcal{G}}(\underline{\theta}_i) = R_{\mathcal{G}}(\overline{\theta}_i)$ for any perturbation of redistributive preferences. In particular, the second part of Appendix Proposition 5 further shows that, under mean-preserving perturbations ($\Delta\tilde{\lambda}_i = 0$), the induced gains and losses must necessarily occur at the type boundaries in the opposite direction from the interior reallocation, whatever the direction of preference changes.

[19] The equivalence between the two equations stems from the observation that the aggregate proportional change in the effective social weight is also the proportional change in the scarcity rent

$$\frac{\Delta\varepsilon_i(s)}{\varepsilon_i(s)} = \mathbb{E}_i^w[R_{\mathcal{G}}(\theta) \mid s]\,.$$

Similarly to the capacity extension expression $\partial_k CS$, the individual surplus change follows from $\mathbb{E}CS_i^{\times}$

$$\Delta\mathbb{E}CS_i(\theta) = \Delta\mathbb{E}\underline{CS}_i + \int_{\underline{\theta}_i}^{\theta} \mathbb{E}_s\left[\Delta q_i(\tilde{\theta}, s) u(q_i(\tilde{\theta}, s), s)\right] d\tilde{\theta}$$

Using lemma 3, the following expression characterizes the change of available surplus $\mathbb{E}\underline{CS}_i$ due to a perturbation of preferences

$$\Delta\mathbb{E}\underline{CS}_i = \mathbb{E}_{(s,i)}\left[\varepsilon_i(s)\, R_\theta(\theta)\, R_q(\theta)\, (R_{\mathcal{G}}(\theta) - \mathbb{E}_i^{w}[R_{\mathcal{G}}(\theta) \mid s])\, \mathbf{1}_{\{s \in T_i\}}\right].$$

Proposition 1 (Surplus and preference perturbations)**.** *Assume Assumption CARA and unimodality of $R_{\mathcal{G}}(\theta)$. Define the set of cutoffs*

$$\tilde{\Theta}_i^{\Delta} := \left\{\tilde{\theta}_i \in [\underline{\theta}_i, \overline{\theta}_i] : \Delta\mathbb{E}CS_i(\tilde{\theta}_i) = 0\right\}.$$

IR-unconstrained allocation. *The pattern of $\Delta\mathbb{E}CS_i(\theta)$ is jointly determined by two signs: $\mathrm{Cov}_i(R_\theta(\theta), R_{\mathcal{G}}(\theta))$, which governs the starting value,*

$$\mathrm{sign}(\Delta\mathbb{E}CS_i(\underline{\theta}_i)) = \mathrm{sign}(\mathrm{Cov}_i(R_\theta(\theta), R_{\mathcal{G}}(\theta))),$$

and $R_{\mathcal{G}}(\theta)$, which governs the interior shape. If $R_{\mathcal{G}}(\theta) \geq 0$ for all θ: expected surplus follows the pattern $+, -, +, -$ with at most three interior cutoffs when $\mathrm{Cov}_i(R_\theta(\theta), R_{\mathcal{G}}(\theta)) > 0$, and the pattern $-, +, -$ with at most two interior cutoffs when $\mathrm{Cov}_i(R_\theta(\theta), R_{\mathcal{G}}(\theta)) < 0$. If $R_{\mathcal{G}}(\theta) \leq 0$ for all θ, the starting value and all interior inequalities are reversed.

IR-constrained allocation. *The lowest type always lies on the boundary: $\underline{\theta}_i \in \tilde{\Theta}_i^{\Delta}$. There exist at most two additional cutoffs $\theta_1 < \theta_2$ in $(\underline{\theta}_i, \overline{\theta}_i)$. If $R_{\mathcal{G}}(\theta) \geq 0$ for all θ, expected surplus decreases for $\theta \in (\underline{\theta}_i, \theta_1)$, increases for $\theta \in (\theta_1, \theta_2)$, and decreases for $\theta \in (\theta_2, \overline{\theta}_i)$. If $R_{\mathcal{G}}(\theta) \leq 0$ for all θ, the inequalities are reversed.*

Proof. See Appendix C.6 □

Proposition 1 delivers a clear redistributive policy result. Even when the direction of a preference shift is unambiguous (as in Lemma 7), the resulting surplus changes are generically non-monotone across types.[20]

In both regimes, the origin of the non-monotonicity in the change in surplus is the same: the individual surplus change combines an IC-rent channel and a level channel, and neither of them follows the direction of the preference perturbation mechanically. Figure 3 provides a simple geometry illustration of the proposition.

The first channel is the IC-rent channel. By construction, informational rents are built from the allocation received by lower types. As a result, when the perturbation changes the quantity assigned to some type $\tilde{\theta}$, this

[20] Assumption CARA is used only to sharpen this result by delivering a closed-form expression for the boundary term $\mathbb{E}\underline{CS}_i$. The unimodality of $R_{\mathcal{G}}(\theta)$ is used only to limit the number of sign changes of $\Delta\mathbb{E}CS_i(\theta)$. If $R_{\mathcal{G}}(\theta)$ is not unimodal, additional crossings may arise, so the proposition should be read as a conservative characterization of the possible alternating surplus regions.

change is inherited by all higher types through the rent formula. The shape of the IC-rent term is therefore determined by the shape of $\Delta q_i(\theta, s)$ across types. Lemma 3 shows that this behavioral response is itself governed by the gap between the proportional change in effective social weight, $R_{\mathcal{G}}(\theta)$, and its equilibrium average. Since $R_{\mathcal{G}}(\theta)$ is non-monotone, the induced quantity response is also generically non-monotone. Hence, even when the underlying preference shift is monotone, the associated IC-rent effect on surplus need not be monotone across types.

The second channel is a level effect, whose source depends on the regime. When the participation constraint is slack, the change in surplus contains both the fixed component $\Delta\mathbb{E}\underline{CS}_i$ and the IC-rent term. The fixed component is common across types and depends on how the perturbation changes the total surplus available for redistribution. In the two-type example, differentiating $\mathbb{E}\underline{CS}_{LH}$ and using Δq_{LH} yields

$$\Delta\mathbb{E}\underline{CS}_i = \frac{\varepsilon}{\alpha} p_L p_H \left(R_\theta(\theta_H) - R_\theta(\theta_L)\right)\left(R_{\mathcal{G}}(\theta_H) - R_{\mathcal{G}}(\theta_L)\right) = \frac{\varepsilon}{\alpha}\mathrm{Cov}(R_\theta(\theta), R_{\mathcal{G}}(\theta)).$$

If a consumer who receives the greatest proportional change also participates most in the revenue, then the available surplus for all consumers increases. When participation binds, the lowest type remains at the boundary, so that $\Delta\mathbb{E}CS_i(\underline{\theta}_i) = 0$. In that case, the level effect is pinned down by the endpoint condition, and the entire surplus profile is shaped by the accumulation of the non-monotone quantity response through the IC formula.

Taken together, these two channels explain why the surplus effect of a redistributive preference perturbation is generically non-monotone across types. A shift in preferences toward higher or lower types does not translate mechanically into a monotone pattern of winners and losers, because the perturbation operates through the endogenous reallocation of quantities and through the way informational rents accumulate along the type space.

I illustrate in Figure 10 in the Appendix a numerical simulation. I assume an initial category characterized by utilitarian preferences, and then compute two (mean-preserving) perturbations in welfare weights $\lambda_i(\theta)$: case A exhibits higher weights on higher types, while case B exhibits higher weights on lower types. In the second panel, $R_{\mathcal{G}}(\theta)$ is positive for case A and negative for case B, following Lemma 7. The third panel shows how the optimal allocation responds to the perturbation: for instance the allocation of $\overline{\theta}_i$ decrease in case A and increase in case B while $\lambda_i(\theta)$ is the highest in case A and the lowest in case B. The fourth panel confirms the illustration in Figure 3.

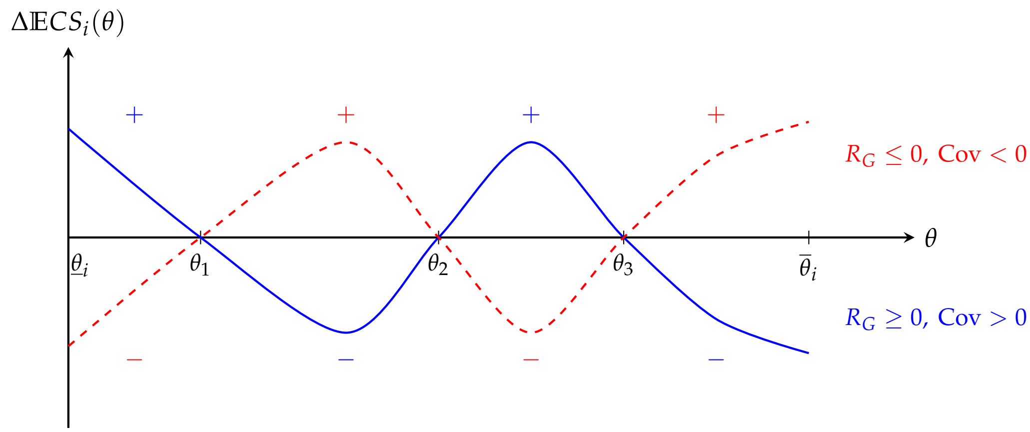

Figure 3: Surplus effect of a preference perturbation across types. Each curve represents the difference of surplus between an initial optimal allocation under a given function of redistributive preferences and a new optimal allocation under a different function of redistributive preferences with respect to types. The blue curve represents a tilt of preferences towards higher types which yields a higher available revenue for the market designer and the red curve represents a tilt of preferences towards lower types which yields a lower available revenue for the market designer.

4 Redistributive Allocation Across Observable Categories

Having characterized the constrained mechanism for a given pair (k_i, I_i), I now turn to the designer's problem of allocating total capacity and total revenue requirements across categories. The previous section provides the objects that govern this outer allocation. The question here is how those same forces translate into differences in the marginal values of capacity and revenue across categories, and hence into the optimal assignment of capacity and budget contributions across observable groups.

4.1 Characterization

Since the objective and category-level constraints are separable across categories, with interaction only through the aggregate constraints on capacity and revenue, the designer's problem decomposes into category-specific value functions $V_i(k_i, I_i)$ and an outer allocation problem over $\{k_i, I_i\}_{i \in N}$. Given the value functions of the single-category problem $V_i(k_i, I_i)$, the designer chooses $\{k_i, I_i\}_{i \in N}$ to solve

$$\max_{\{k_i, I_i\}_{i \in N}} \sum_i \mu_i \, V_i(k_i, I_i)$$

subject to

$$\sum_i \mu_i k_i = k, \qquad \sum_i \mu_i I_i = I(k), \qquad k_i \geq 0, \quad I_i \geq 0 \quad \forall i \in N.$$

where $V_i(k_i, I_i)$ is defined as

$$V_i(k_i, I_i) \ := \ \max_{q_i(\theta,s)} \ \mathbb{E}_{(s,i)}[\Gamma_i(\theta)\, U(q_i(\theta,s),s)] - \tilde{\lambda}_i I_i$$

subject to incentive compatibility and the category-specific constraints. By the envelope theorem, the two partial derivatives of V_i are

$$\frac{\partial V_i(k_i, I_i)}{\partial k_i} \ = \ \bar{\varepsilon}_i \ , \qquad \frac{\partial V_i(k_i, I_i)}{\partial I_i} \ = \ -\tilde{\lambda}_i - \beta_i(k_i, I_i).$$

where $\bar{\varepsilon}_i = \mathbb{E}_s \left[\varepsilon_i(s)\, \mathbf{1}_{\{s \in T_i\}}\right]$.

The envelope decomposition shows that the problem is governed by only two marginal objects. First,

$$\Pi_i(k_i, I_i) \ := \ \frac{\partial V_i(k_i, I_i)}{\partial k_i} \ = \ \mathbb{E}_s \left[\varepsilon_i(s)\, \mathbf{1}_{\{s \in T_i\}}\right] \ = \ \mathbb{E}_{(s,i)} \left[u(q_i(\theta,s),s)\, \mathcal{G}_i(\theta)\, \mathbf{1}_{\{s \in T_i\}}\right],$$

is the expected shadow value of relaxing the category-specific capacity constraint. As in canonical peak-load pricing, the marginal value of an additional unit of capacity is equal to the expected scarcity rent. Equivalently, using the first-order condition FOC_q, it can be written as the expected aggregate marginal utility generated by extra capacity. The difference relative to the canonical model is that this marginal utility is here weighted by the effective social weight $\mathcal{G}_i(\theta)$, rather than by θ. This IC-adjusted weighting will be crucial for the comparative statics below.[21]

Second,

$$\frac{\partial V_i(k_i, I_i)}{\partial I_i} = -\tilde{\lambda}_i - \beta_i(k_i, I_i),$$

is the effective marginal social cost of requiring one additional unit of revenue from category i. The term β_i captures the canonical marginal cost of tightening the budget, while $\tilde{\lambda}_i$ captures the additional social cost induced by redistributive concerns.

Lemma 4 (Across-category allocation)**.** *The outer problem admits at least one solution, and the set of optimal allocations is convex. The optimality conditions are, for each $i \in N$,*

$$\Pi_i(k_i^*, I_i^*) \ = \ \zeta_k, \qquad \tilde{\lambda}_i + \beta_i(k_i^*, I_i^*) \ = \ -\zeta_I + \xi_i,$$

with complementary slackness $\xi_i I_i^* = 0$ for all i and where $\zeta_k, \zeta_I \in \mathbb{R}$ denote the multipliers on the capacity and revenue equality constraints, and $\xi_i \geq 0$ the multiplier on $I_i \geq 0$.[22]

Proof. See Appendix C.7 □

[21] Note that $\Pi_i(k_i, I_i)$ is strictly positive whenever category i is capacity-constrained with positive probability, since $\varepsilon_i(s) > 0$ in every on-peak state.

[22] The Inada conditions on $u(\cdot)$ ensure that $k_i > 0$ at the optimum, so the multiplier on $k_i \geq 0$ is zero.

The two optimality conditions have the usual economic interpretation. The designer allocates capacity until the marginal value of every categories, measure by $\Pi_i(k_i^*, I_i^*)$, is the same across all categories. He equalizes the effective marginal social cost of revenue extraction $\tilde{\lambda}_i + \beta_i(k_i^*, I_i^*)$ across categories that contribute to the budget.

4.2 Ranking of the optimal allocation

I now study how redistributive preferences shape the allocation of capacity and revenue across categories. I study two rankings: the capacity ranking, which determines how much capacity each category receives, and the revenue ranking, which determines how much each category contributes to the budget, the latter decomposing into an extensive margin governing which categories contribute at all, and an intensive margin governing how much each contributing category pays. I begin with a local comparison between two nearby categories, representing category j as a marginal variation of category i, and then use this comparison to derive a global ordering along a path of redistributive preferences. Throughout, all categories differ in their welfare weight functions $\lambda_i(\theta)$ and share the same support Θ.

Capacity ranking. Using the first-order conditions in Lemma 4, consider the induced difference in the marginal value of capacity around category i:

$$\Delta\Pi_i(k_i^*, I_i^*) = \mathbb{E}_{(s,i)}[(\underbrace{u_q(q_i(\theta,s),s)\,\mathcal{G}_i(\theta)\,\Delta q_i(\theta,s)}_{\text{behavioral effect}} + \underbrace{u(q_i(\theta,s),s)\,\Delta\mathcal{G}_i(\theta)}_{\text{weight effect}})\mathbf{1}_{\{s\in T_i\}}].$$

This expression separates the effect of a preference variation into two margins: a behavioral effect, through the induced change in the optimal allocation $q_i(\theta,s)$, and a direct reweighting effect, through the change in effective social weights $\Delta\mathcal{G}_i(\theta)$. By Lemma 3 and from Δq, the allocation response has the same decomposition. At the aggregate level, however, the two direct reweighting terms offset each other. Therefore, the change in the marginal value of capacity is equal to the average proportional difference in effective social weights $\mathbb{E}_i^w[R_{\mathcal{G}}(\theta) \mid s]$.

Theorem 2 (Capacity ordering)**.** *Suppose category j differs from a baseline category i only through a marginal perturbation of redistributive preferences. Fix the baseline optimum (k_i^*, I_i^*). Then, to first order in $\Delta\lambda$, the following capacity ranking holds*

$$k_j^* > k_i^* \quad \Longleftrightarrow \quad \mathbb{E}_s\Big[\varepsilon_i(s)\,\mathbb{E}_i^w[R_{\mathcal{G}}(\theta) \mid s]\,\mathbf{1}_{\{s\in T_i\}}\Big] > 0.$$

Under Assumption CARA, the condition reduces to $k_j^* > k_i^*$ if and only if $\mathbb{E}_i[R_{\mathcal{G}}(\theta)] > 0$.

Proof. See Appendix C.8 □

The theorem provides a direct outer-margin counterpart to the single-category analysis. A category receives more capacity when an additional unit of capacity has a higher marginal social value, and the theorem shows that this marginal value is governed by the weighted average of the proportional difference in effective social weights. The same object $R_{\mathcal{G}}(\theta)$ that governed the reallocation of quantities and the surplus effects of preference differences within a category now determines which category receives more capacity at the outer optimum. Therefore, the relevant comparison is not which category has the larger average effective weight nor the absolute size of the difference in redistributive preferences, but whether category j's advantage in effective weights is concentrated on types for which category i's effective weight is low.

I provide, in Figure 11 in the Appendix, simple numerical examples of redistributive preferences around the utilitarian benchmark, that is, $\lambda_i(\theta) = 1$. These examples show that the ranking induced by the absolute effect, $\mathbb{E}_i[\Delta\mathcal{G}_i(\theta)]$, and the ranking induced by the proportional effect, $\mathbb{E}_i[R_{\mathcal{G},i}(\theta)]$, can be non-trivial and may even go in opposite directions. In two cases, I consider pairs of categories i and j whose redistributive weights $\lambda_i(\theta)$ and $\lambda_j(\theta)$ both deviate from the constant benchmark 1, while all other parameters are held fixed. I show that even when category i has a higher average effective weight relative to the utilitarian benchmark whereas category j has a lower one, so that $\mathbb{E}_i[\Delta\mathcal{G}_i(\theta)] > 0 > \mathbb{E}_j\left[\Delta\mathcal{G}_j(\theta)\right]$, it is still possible to have $\mathbb{E}_i[R_{\mathcal{G},i}(\theta)] < 0 < \mathbb{E}_j\left[R_{\mathcal{G},j}(\theta)\right]$.

I now derive global implications of Theorem 2 by imposing additional structure on redistributive preferences. The theorem provides a local comparison between nearby categories; the role of the next assumption is to place all categories along a common ordered path, so that these local comparisons can be aggregated into a global ranking.

Assumption Order (Ordered preference path)**.** *For each adjacent pair $i-1$ and i, there exists a positive C^1 path $(t,\theta) \mapsto \lambda(t,\theta)$ on $[0,1] \times \Theta$ such that $\lambda(0,\theta) = \lambda_{i-1}(\theta)$ and $\lambda(1,\theta) = \lambda_i(\theta)$. Along this path, two conditions hold for every $t \in [0,1]$:*

(i) the average redistributive weight is weakly increasing, $\partial_t\tilde{\lambda}(t) \geq 0$, where $\tilde{\lambda}(t) := \int_\Theta \lambda(t,\theta)dG(\theta)$,

(ii) the relative tilt of redistributive weights is toward higher types, in the sense that $\partial_t \log\lambda(t,\theta)$ is weakly increasing in θ on Θ.

Assumption Order orders categories along a common redistributive path. Condition (i) requires average welfare weights to be weakly increasing along the ordering. Condition (ii) requires this increase to be proportionally stronger for higher types. Together, these two restrictions allow the local comparison in Theorem 2 to be iterated across adjacent categories. By Lemma 7 in Appendix A.4, they imply a monotone ordering of effective social weights that carries over to the equilibrium capacity profile.

Corollary 1 (Capacity ordering along an ordered preference path)**.** *Consider an equilibrium allocation* $\{k_i^*, I_i^*, \beta_i^*\}_{i \in N}$. *Under Assumption Order, the equilibrium capacity profile is globally weakly monotone:*

$$k_1^* \leq \cdots \leq k_n^*.$$

Proof. See Appendix C.10 □

The corollary provides a sufficient condition for a global ranking of categories by capacity allocation. Capacity depends on both the average welfare weight $\tilde{\lambda}_i$ and the distribution of redistributive preferences across types. Categories that place proportionally more weight on higher types have a larger $\mathcal{G}_i(\theta)$, not only through the average term $\tilde{\lambda}_i + \beta_i$, but also through the informational-rent component $\Lambda_i(\theta)$. Under Assumption Order, these two forces move in the same direction. As a result, categories that are higher in the redistributive ordering optimally receive a larger capacity.

In Figure 12 in the Appendix, I provide a numerical illustration of the sufficient condition. I compute the optimal allocation of k_i across 50 categories for a given aggregate capacity k. I consider two cases. In both, categories are ordered so that, between two adjacent categories i and $i+1$, the average welfare weight is strictly increasing. In the first case, redistributive preferences are gradually tilted toward higher types, whereas in the second they are gradually tilted toward lower types. For comparability, I assume that, for a given category i, the average welfare weight is the same across the two cases. I show that, in the second case, capacity may be increasing across some initial categories, but that the ranking eventually reverses. In the first case, which satisfies the condition of the corollary, the optimal allocation k_i^* is clearly strictly increasing.

Revenue ranking. To rank categories in terms of revenue requirements, I again use the first-order condition from Lemma 4. For a local comparison around category i, a first-order expansion of the outer condition for category j gives

$$0 = \Delta\tilde{\lambda} + \frac{\partial \beta_i}{\partial k_i}(k_i^*, I_i^*)\,\Delta k + \frac{\partial \beta_i}{\partial I_i}(k_i^*, I_i^*)\,\Delta I + \Delta_\lambda \beta_i(k_i^*, I_i^*),$$

where $\Delta_\lambda \beta_i(k_i^*, I_i^*)$ denotes the direct effect of the preference perturbation on the category-specific shadow cost at fixed (k_i^*, I_i^*), obtained from the implicit-function derivative of the binding participation condition while holding β_i fixed in the numerator. This expression shows that the comparison of equilibrium revenue requirements across categories depends on two forces: the difference in average welfare weights, $\Delta\tilde{\lambda}$, and the endogenous adjustment of the shadow cost β_i, which reflects both the induced change in capacity and the direct effect of the preference perturbation. The single-category analysis provides the sign of $\frac{\partial \beta_i}{\partial k_i}$ and characterizes $\Delta_\lambda \beta_i$, yielding the following result.

Proposition 2 (Revenue ordering)**.** *Consider an equilibrium allocation $\{k_i^*, I_i^*\}_{i\in N}$. Define the two endogenous subsets*

$$N^+ := \{i \in N : \beta_i^* > 0\}, \qquad N^0 := \{i \in N : \beta_i^* = 0\},$$

which partition N into categories that contribute to the budget and categories that do not contribute to the budget (either may be empty). Then:

Extensive margin. *At any equilibrium, $\max_{i\in N^+} \tilde{\lambda}_i < \min_{i\in N^0} \tilde{\lambda}_i$, so N^+ and N^0 are strictly separated in terms of average welfare weights, regardless of the within-type structure of preferences.*

Intensive margin. *Fix a baseline category $i \in N^+$, and let category $j \in N^+$ be obtained from i through a marginal perturbation of redistributive preferences with $\Delta\tilde{\lambda} := \tilde{\lambda}_j - \tilde{\lambda}_i$ and $\Delta k := k_j^* - k_i^* > 0$. Then to the first order*

$$I_j^* < I_i^* \quad \Longleftrightarrow \quad \underbrace{-\left(\Delta k \frac{\partial \beta_i}{\partial k_i}(k_i^*, I_i^*) + \Delta_\lambda \beta_i(k_i^*, I_i^*)\right)}_{\text{revenue effect}} < \underbrace{\Delta\tilde{\lambda}}_{\text{preference effect}} .$$

where the sign of $\Delta k \frac{\partial \beta_i}{\partial k_i}(k_i^*, I_i^*)$ is determined by the sign of EuJ and $\Delta_\lambda \beta_i(k_i^*, I_i^*)$ is the first-order effect of the preference perturbation on β_i at fixed (k_i^*, I_i^*). Under Assumption CARA, and using the first-order expression for Δk, this reduces to

$$I_j^* < I_i^* \quad \Longleftrightarrow \quad \frac{\mathbb{E}_i[R_{\mathcal{G}}(\theta)]\,\mathbb{E}_i[R_\theta(\theta)]}{\mathbb{E}_i[R_\theta(\theta)^2]} + \frac{\mathrm{Cov}_i(R_\theta(\theta), R_{\mathcal{G}}(\theta))}{\mathrm{Var}_i(R_\theta(\theta))} < 0.$$

Proof. See Appendix C.11 □

The proposition completes the characterization on the revenue side. The extensive margin is governed solely by average welfare weights: categories with sufficiently low average redistributive weight contribute to the budget, whereas categories with sufficiently high average redistributive weight do not. This margin compares a common shadow value of public funds, $\tilde{\lambda}^*$, to the category-specific marginal social cost of raising revenue. Category i contributes only when the social value of one more unit of revenue exceeds both the redistributive cost of taking it from the category, $\tilde{\lambda}_i$, and the additional incentive distortion it creates, β_i^*. I do not rule out the presence of a marginal category that contributes positively to the budget, even though its budget constraint is slack at the optimum. When such a category exists, however, it is necessarily unique.

Within the contributing set, however, equilibrium revenue requirements depend not only on average welfare weights but also on the category's ability to raise revenue, through both the induced change in equilibrium capacity and the direct effect of the preference perturbation on the shadow cost of revenue. Under Assumption CARA, this tradeoff becomes especially transparent, because the revenue-ordering condition can be expressed entirely in terms of the same sufficient statistics that govern the previous sorting results, and the effect of the difference in average welfare weights cancels out. The capacity difference is summarized by the marginal capacity value and captured by the proportional change $\mathbb{E}_i[R_{\mathcal{G}}(\theta)]$ (Theorem 2). The effect of a larger capacity on the shadow cost of revenue is summarized by the distribution of virtual surplus across types, $\mathbb{E}_i[R_\theta(\theta)]$ (Theorem 1). Finally, the direct effect of the preference perturbation on the ability to raise revenue is captured by the covariance term, as in Proposition 1. Under CARA, the revenue ranking is therefore entirely governed by the sorting of consumers along these two objects, independent of the sign of $\Delta\tilde{\lambda}$.

5 Optimal investment

The third and final step endogenizes total investment. I now let k itself be a choice variable. The designer solves

$$\max_{k\geq 0} W(k),$$

where $W(k)$ denotes the outer value function once total revenue is tied to the endogenous investment cost $I(k)$, with I strictly increasing and convex.

Increasing k relaxes scarcity in states where capacity binds, but it also requires additional revenue to finance the larger investment. The optimal investment level is therefore pinned down by equating the marginal value of relaxing the capacity constraint with the marginal financing cost of additional capacity. As shown in Lemma 5, this condition is governed by the same objects that organize the previous analysis.

Lemma 5 (Optimal investment)**.** *The investment problem* $\max_{k\geq 0} W(k)$ *admits a unique interior solution* $k^* > 0$*. The optimal investment level satisfies*

$$I'(k^*) = \frac{1}{\mu^+\tilde{\lambda}^*}\sum_{i\in N^+}\mu_i\,\Pi_i(k_i^*, I_i^*) = \frac{1}{\mu^+}\sum_{i\in N^+}\mu_i\,\mathbb{E}_{(s,i)}\left[u(q_i(\theta,s),s)\,\frac{\mathcal{G}_i(\theta)}{\tilde{\lambda}^*}\,\mathbf{1}_{\{s\in T_i\}}\right],$$

where $\mu^+ := \sum_{i\in N^+}\mu_i$ *denotes the mass of contributing categories and* $\tilde{\lambda}^* := -\zeta_I > 0$ *is the equilibrium marginal cost of funds. Moreover,*

$$\max_{i\in N^+}\tilde{\lambda}_i < \tilde{\lambda}^* \leq \min_{i\in N^0}\tilde{\lambda}_i.$$

Proof. See Appendix C.12 □

Lemma 5 is the incomplete-information counterpart of the classic peak-load investment rule. In the complete-information benchmark with redistributive preferences, incentive constraints play no role. There are no informational rents, so $J_i(\theta) = \theta$ and $\Lambda_i(\theta) = 0$. Hence the normalized incomplete-information object collapses to its complete-information counterpart:

$$\frac{\mathcal{G}_i(\theta)}{\tilde{\lambda}^*} \longrightarrow \theta,$$

This contrasts with the utilitarian incomplete-information benchmark. There, private information prevents the spot market from implementing the optimum at a given capacity, but this distortion disappears at the optimal investment level in the canonical model. Intuitively, the optimal investment condition pins down the same object that governs the redistributive incidence of a capacity expansion, so that at k^* the mechanism lies in the IR-unconstrained regime, which is implementable by the spot market (Spulber, 1992b; Chao and Wilson, 1987). Under redistributive preferences, by contrast, the investment condition is no longer governed by that same object. As informational rents are no longer welfare-neutral, the restoration result fails, and a uniform spot price does not implement the constrained-optimal allocation even at k^*.

Figure 13 in the Appendix provides a numerical illustration of welfare $W(k)$ under four scenarios: utilitarian versus redistributive preferences, and complete versus incomplete information, holding the primitives fixed. The figure shows that, under redistributive preferences, the complete- and incomplete-information welfare profiles differ even at the optimum. By contrast, under utilitarian preferences, the two profiles coincide at the optimal investment level k^*.

I conclude by using the capacity- and preference-sorting margins developed in the previous analysis to characterize how redistributive preference perturbations affect both the optimal investment level and the incidence of a global capacity expansion across categories.

Proposition 3 (Investment and preference perturbations)**.** *Assume Assumption CARA. Consider a mean-preserving local perturbation of redistributive preferences, evaluated at the baseline equilibrium $(k^*, \{k_i^*, I_i^*\}_{i\in N})$. Let $\Delta\tilde{\lambda}^*$ denote the resulting direct effect on the common marginal cost of funds. Then*

$$W''(k^*)\,\Delta k^* = \underbrace{-\sum_i \mu_i\,\Pi_i\,\mathbb{E}_i[R_{\mathcal{G},i}]}_{\text{average effect}} + \underbrace{I'(k^*)\,\Delta\tilde{\lambda}^*}_{\text{shadow-cost effect}},$$

where

$$\Delta\tilde{\lambda}^* = -\frac{\sum_{i\in N^+} \mu_i\,\bar{\varepsilon}_i\,\mathrm{Cov}_i(R_{\theta,i}, R_{\mathcal{G},i})}{\sum_{i\in N^+} \mu_i\,\bar{\varepsilon}_i\,\mathrm{Var}_i(R_{\theta,i})}.$$

Here $W''(k^*) < 0$ and $\bar{\varepsilon}_i \geq 0$.

Proof. See Appendix C.13. □

The proposition shows that redistributive preference perturbations affect optimal investment through the same two forces as in the previous analysis. The first operates through the preference-sorting margin, by changing the category-level marginal value of capacity Π_i as allocations are reshaped through $R_{\mathcal{G},i}$. The second operates through the capacity-sorting margin, by changing the equilibrium marginal cost of funds $\tilde{\lambda}^*$, with the sign again governed by the covariance between the redistributive shift and the type profile. At the aggregate level, these two effects jointly determine whether optimal investment rises or falls.

I illustrate this result in Figure 14 in the Appendix, where I simulate how the optimal investment level k^* responds to redistributive preference changes. I start from a benchmark with two utilitarian categories, i and $i+1$. I then iteratively perturb preferences and recompute the optimal investment level for each successive pair of categories: in the first step, category i remains utilitarian while category $i+1$ is a marginal perturbation of i; in the next step, category $i+1$ becomes the new baseline and category $i+2$ is constructed as a perturbation of it; and so on. I consider two cases. In the first, redistributive preferences are increasingly tilted toward higher types; in the second, they are increasingly tilted toward lower types. In the numerical examples I consider, the sorting and average effects move in the same direction under monotone mean-preserving perturbations, producing a clear opposite shift in optimal investment.

Finally, I connect a marginal increase in aggregate capacity k to its redistributive incidence across categories and types. While I do not provide a formal characterization here, the previous results already identify the two objects that govern this effect: the incidence on a single category of k_i, and the response of equilibrium category capacities k_i^* to a change in aggregate capacity k. This yields a simple condition that captures the main economic mechanism at work. Note that the response of the equilibrium revenue I_i relies on the same mechanism. Consider a neighborhood of an interior optimum in which the set of active and contributing categories remains unchanged. Then

$$\frac{d\,\mathbb{E}CS_i(\theta)}{dk} = \frac{\partial \mathbb{E}CS_i(\theta)}{\partial k_i}\,\frac{dk_i^*}{dk}.$$

The first term is the single-category incidence characterized in Theorem 1. The second term captures how a marginal increase in aggregate capacity is redistributed across categories. As shown in Appendix Lemma 8, this term depends on two objects. The first is the category-specific expression

$$\Omega_i^{-1}\left(\mathbb{E}^{w}_{(s,i)}[uJ \mid s] - I'(k^*)\right), \quad \text{where} \quad \Omega_i := \mathbb{E}_s\left[\frac{1}{\overline{w}_i(s)}\mathbf{1}_{\{s\in T_i\}}\right].$$

The first term measures whether the marginal revenue EuJ for category i is high or low relative to the marginal investment cost. The second is the corresponding aggregate expression,

$$\sum_{j\in N}\mu_j\,\Omega_j^{-1}\left(\mathbb{E}^{w}_{(s,j)}[uJ \mid s] - I'(k^*)\right), \qquad (d\lambda^*)$$

which governs how the equilibrium marginal cost of funds $\tilde{\lambda}^*$ responds to a change in aggregate capacity.[23]

This yields a simple interpretation of the redistributive effect of a global capacity expansion. The allocation of additional capacity depends on both the category-level comparison between EuJ and $I'(k^*)$, and the corresponding aggregate object $d\lambda^*$. The sign of $d\lambda^*$ determines the direction of the reallocation: if $d\lambda^*$

[23]In this sense, the aggregate response mirrors the previous proposition on $\Delta\tilde{\lambda}^*$: what matters, again, is how the marginal revenue compares with the relevant aggregate benchmark.

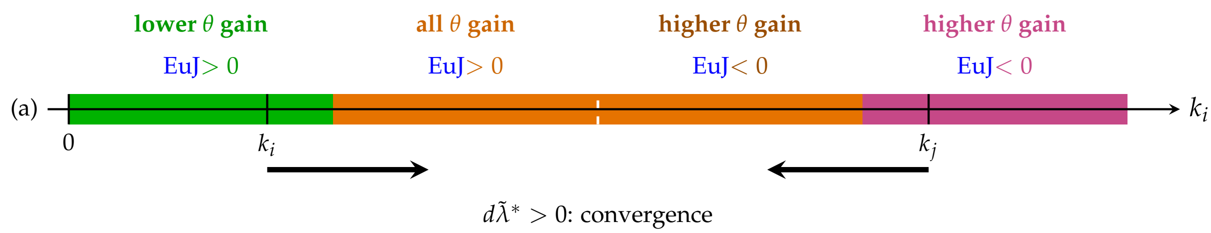

Figure 4: Redistributive incidence of a global capacity expansion for two categories i and j, initially located in the lower-θ-gain and higher-θ-gain regimes, respectively. The example assule $d\tilde{\lambda}^* > 0$, so the expansion reallocates capacity toward k_i and away from k_j, compressing redistributive heterogeneity across categories.

is positive, categories with relatively high marginal revenue receive more of the expansion; if $d\lambda^*$ is negative, categories with relatively low marginal revenue receive more. Moreover, as shown in the section A.3 of the Appendix and used in Theorem 1, EuJ moves monotonically with k_i: depending on $R_\theta(\theta)$ and $R_q(\theta)$, it may switch from positive to negative, or from negative to positive. This yields a qualitative prediction based on these two ratios and the sign of $d\lambda^*$.

To illustrate, take two categories i and j, and assume $R'_\theta(\theta) > 0$, so that in both categories the virtual term $u(q_i(\theta, s), s)J_i(\theta)$ is increasing in θ. Suppose that, at the initial allocation, $k_j^* \gg k_i^*$, and that EuJ is positive for the low-capacity category i but negative for the high-capacity category j. If $d\lambda^*$ is positive, the global expansion reallocates relatively more capacity toward category i, compressing category-level allocations and generating convergence. If $d\lambda^*$ is negative, the opposite occurs: the expansion reallocates relatively more capacity toward category j, widening category-level allocations and generating divergence. Figure 4 illustrates this mechanism.

6 Application to Nonlinear Tariffs

The optimal allocation characterized above equates the scarcity rent to each type's marginal utility weighted by its effective social weight. As a result, a tariff with a single variable price, such as spot pricing, generally fails to implement the optimum, even when combined with type-specific lump-sum transfers. In network industries such as electricity and water, nonlinear tariffs, and in particular block tariffs, are often used for screening and redistributive purposes, but their optimal structure in a peak-load environment with redistributive preferences and incentive constraints remains unknown. This section uses the theoretical framework developed above to characterize the optimal nonlinear tariff and to provide an economic reading of the mechanism.

I begin with the single-category case and examine how the shape of the marginal price schedule depends on the distribution of redistributive preferences across types. I then compare the tariff menus across observable categories when tagging is allowed, and contrast the resulting individual surpluses with two benchmarks: a uniform spot price and a no-discrimination policy that pools all categories under a common tariff.[24] Finally, I examine how a capacity expansion shifts the marginal price schedule across types and categories, and relate these menu changes to the redistributive incidence results established in the main text.

Following the nonlinear price literature (Spence, 1977; Maskin and Riley, 1984; Spulber, 1992a), the same allocation can be represented by a tariff with a fixed component, equal to the available surplus $\mathbb{E}\underline{CS}_i$ (if it exists), and a variable component obtained by integrating the marginal price over quantity. Let $\theta^*(q,s)$ denote the unique type satisfying $q_i(\theta^*(q,s),s) = q$, which is well defined since $q_i(\theta,s)$ is strictly increasing in θ on T_i under Assumption Int. The marginal price at quantity q in state $s \in T_i$ is the willingness to pay of the type who chooses that quantity at the optimum:

$$p_i(q,s) := \theta^*(q,s)\, u(q,s) = \theta^*(q,s)\, \frac{\varepsilon_i(s)}{\mathcal{G}_i(\theta^*(q,s))}.$$

Off-peak ($s \in S_i$), the marginal price is zero. Evaluated along the optimal allocation $q = q_i(\theta,s)$, so that $\theta^*(q_i(\theta,s),s) = \theta$, this becomes

$$p_i(q_i(\theta,s),s) = \theta\, \frac{\varepsilon_i(s)}{\mathcal{G}_i(\theta)}.$$

A block tariff is then obtained by discretizing this nonlinear schedule over a partition of the quantity space and assigning to each block the average marginal price over the corresponding interval. As the partition becomes finer, the block tariff converges to the underlying nonlinear tariff. Throughout the discussion, the results are derived under a CRRA utility specification, a uniform distribution of types, and exponentially tilted welfare weights. Specifically,

$$U(q,s) = (1+s)\frac{q^{0.6}}{0.6} - q, \qquad \theta \sim \mathcal{U}\,[0.5, 1.5]\,, \qquad \lambda_i(\theta) = \tilde{\lambda}_i\, \frac{\exp\left(a_i(\theta-1)\right)}{\int_{0.5}^{1.5} \exp\left(a_i(t-1)\right) g_i(t)\, dt}.$$

Where a_i is the slope of the welfare weights. If it is positive, the designer favors higher types; if it is negative, the designer favors lower types. With this type distribution, all types below 0.75 have a negative virtual surplus $J_i(\theta)$.

I begin with the single-category case. Fix a category i and an allocation (k_i, I_i), and compare the shape of the implemented marginal price schedule under alternative redistributive preferences. A preference perturbation at fixed capacity yields

$$\frac{\Delta p_i(q,s)}{p_i(q,s)} = \mathbb{E}_i^w[R_{\mathcal{G},i} \mid s] - R_{\mathcal{G},i}(\theta) = -\frac{\Delta q_i(\theta,s)}{R_q(\theta)}. \qquad (\Delta_1 p_i)$$

The change in the tariff therefore mirrors the preference-sorting margin from Section 3.5. Therefore, following Lemma 3 and Appendix Proposition 5, for categories that generate available surplus, i.e., $\beta_i = 0$, and more generally for categories with a low revenue requirement, the optimal marginal price schedule is always

[24] The no-discrimination policy is obtained by solving for the optimal allocation by aggregating all types in a single category.

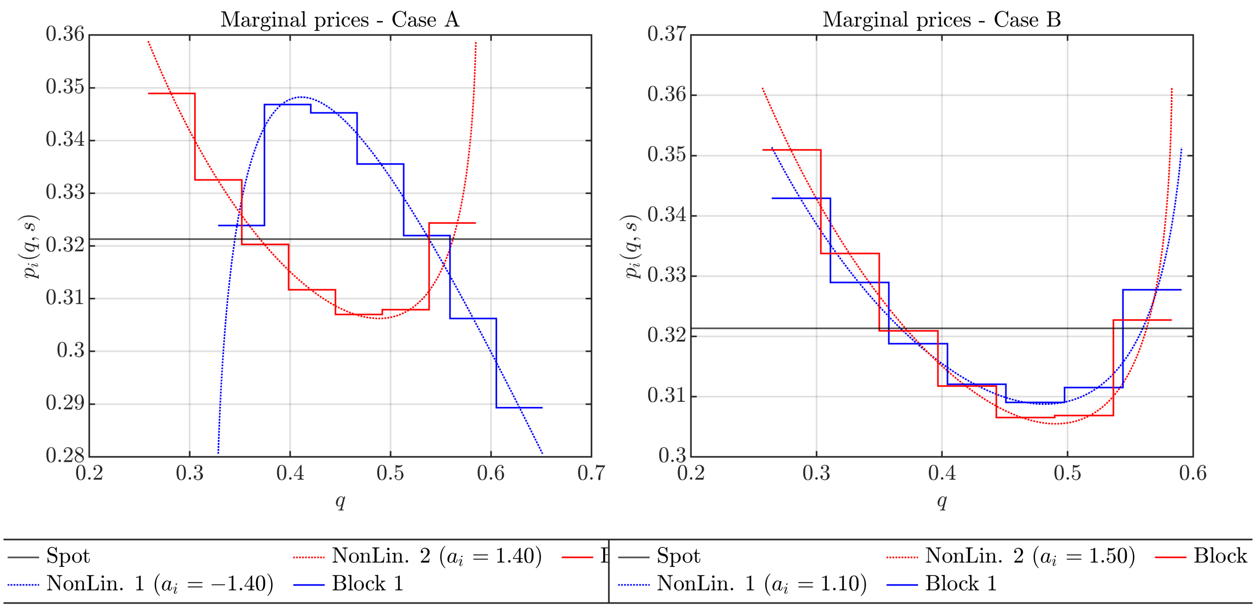

Figure 5: Optimal nonlinear and block-tariff marginal price schedules in the single-category illustration.

non-monotone across types. A direct policy implication is that, for those categories, the optimal marginal price schedule is neither the standard increasing block tariff (IBT) nor the decreasing block tariff (DBT). This stands in clear contrast with the policy view that favored consumers should face an IBT (Borenstein, 2012), since the categories for which redistribution is most salient are precisely those whose optimal menu departs most from standard block tariff formats

Figure 5 represents the optimal marginal price for the nonlinear tariff and the price of the corresponding block tariff.[25] Case A shows that opposite preference directions produce schedules that cross the spot price in opposite directions, reflecting the sign reversal of $R_{\mathcal{G},i}(\theta) - \mathbb{E}_i^w[R_{\mathcal{G},i}|s]$ across types. Similarly, the ambiguity in the sign extends to Case B between two categories with a same direction of perturbations.

I now turn to a local comparison of marginal price schedules across categories. Consider two nearby categories, i and j, that differ in their redistributive weights, capacity allocations, and revenue requirements, as induced by the optimal allocation of aggregate capacity k and aggregate revenue I. A first-order comparison of the implemented marginal prices can be written as

$$\frac{\Delta p_i(q,s)}{p_i(q,s)} = \underbrace{\mathbb{E}_i^w[R_{\mathcal{G}}(\theta) \mid s] - R_{\mathcal{G}}(\theta)}_{\text{preference sorting}} - \underbrace{\Delta k \cdot \frac{\partial q_i(\theta,s)}{\partial k_i} \frac{1}{R_q(\theta)}}_{\text{capacity ranking}} + \underbrace{\Delta I \cdot \frac{\mathbb{E}_i^w[R_\theta \mid s] - R_\theta(\theta)}{\mathbb{E}_{(s,i)}^w[\mathrm{Var}_i(uJ)]}}_{\text{revenue ranking}}, \quad (\Delta_2 p_i)$$

[25]Figure 17 in the Appendix reports the redistributive weight profiles used in the numerical illustration, each constructed as a mean-preserving perturbation of the utilitarian benchmark. Figure 18 in the Appendix shows that, for a fixed perturbation favoring lower types, the standard IBT pattern emerges only when the category faces a sufficiently high revenue requirement, and hence a high β_i.

where $\mathbb{E}^{w}_{(s,i)}[\text{Var}_i(uJ)]$ is a weighted variance term formally defined in 4 in the Appendix, and Δk and ΔI represent the difference in capacity and revenue requirement between two categories. The first term is the single-category expression that captures preference comparison from $\Delta_1 p_i$. The second term captures the capacity ranking effect. The sign of Δk is given by Theorem 2, while the sign of the derivative follows Theorem 1 and is expressed in $\partial_k q_i$. The third term is the intensive margin of revenue ranking from Proposition 2, and also depends on the monotonicity direction of $R_\theta(\theta)$.

I illustrate this decomposition through a configuration of different categories in Figure 6. I considers five tagged categories with distinct average welfare weights: categories 1 and 2 (red) have the same higher average weight $\tilde{\lambda}_i = 1.01$ and do not contribute to revenue; categories 4 and 5 (blue) have the same lower average weight $\tilde{\lambda}_i = 0.99$ and are contributors; and category 3 is the marginal category with $\tilde{\lambda}_3 = \tilde{\lambda}^* = 1$. The Appendix collects the supporting numerical detail, Figure 20 reports the redistributive weight profiles, Figure 23 the capacity and revenue ordering underlying the marginal price ranking, and Figure 21 the three-channel decomposition of $\frac{\Delta p_i(q,s)}{p_i(q,s)}$. As established in Corollary 1, average welfare weights are not sufficient to rank categories by capacity allocation, and the ordering of marginal price schedules inherits the same indeterminacy through the scarcity rent channel. In the configuration shown, categories with higher average welfare weights place relatively more weight on lower types, so the sufficient condition of Corollary 1 is not satisfied and the capacity ranking is no longer determined by average weights alone. As a result, despite having higher average welfare weights, these categories receive less capacity in the example. Their scarcity rents are therefore higher, which pushes their marginal prices above those of categories with lower average welfare weights.[26]

The consumer surplus comparison in Figure 6 clarifies who gains and who loses from the transition from spot pricing to nonlinear tariffs, with and without tagging. In all three configurations, the no-tagging allocation strictly improves upon the spot benchmark for every category, so replacing uniform spot pricing with a common nonlinear tariff is Pareto-improving across categories. By contrast, the move from spot pricing to tagged nonlinear tariffs is not Pareto-improving for every category. Tagging allows the designer to differentiate menus across observable categories and thereby to reallocate surplus across groups. In the three configurations considered here, the categories that lose relative to the spot benchmark are precisely the contributing categories, that is, those with lower average welfare weights. Conversely, the favored categories achieve a higher overall surplus, even though their variable prices may exceed those of less favored categories.

I finally turn to the effect of a global capacity expansion on the implemented nonlinear tariff. I study how the marginal price schedule evolves with total capacity within each category, and compare these tariff adjustments across categories along the equilibrium path. Similar to the previous expressions, the

[26]The Appendix reports two complementary cases. When higher-weight categories instead place relatively more weight on higher types, the capacity and menu ordering align with Corollary 1. When the tilt toward lower types is sufficiently asymmetric and accompanied by a much larger average weight, some marginal prices of favored categories do fall below those of less favored ones.

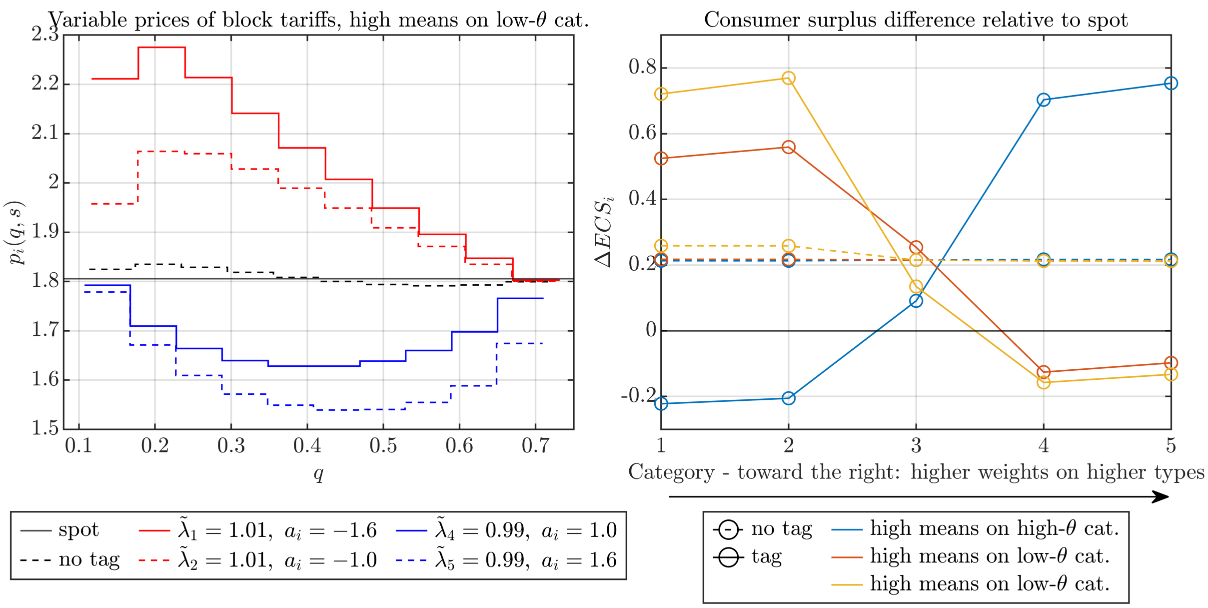

Figure 6: Marginal price schedules by categories and the associated consumer-surplus differences relative to the spot benchmark, with comparison to the no-tagging allocation.

proportional change in the marginal price can be written as

$$\frac{\Delta p_i(q,s)}{p_i(q,s)} = dk\left(\underbrace{\frac{dk_i^*}{dk}\cdot\frac{\partial q_i(\theta,s)}{\partial k_i}\frac{1}{R_q(\theta)}}_{\text{capacity ranking}} + \underbrace{\frac{dI_i^*}{dk}\cdot\frac{\mathbb{E}_i^w[R_\theta \mid s]-R_\theta(\theta)}{\mathbb{E}_{(s,i)}^w[\mathrm{Var}_i(uJ)]}}_{\text{revenue ranking}}\right). \qquad (\Delta_3 p_i)$$

Thus, a global capacity expansion affects the nonlinear tariff through the same margins that govern the allocation problem in $\Delta_2 p_i$: the reallocation of capacity across categories and the induced reallocation of revenue requirements. The implication of $\Delta_3 p_i$ for a non-contributing category, for which $\frac{dI_i^*}{dk} = 0$, is immediate. As shown in the proof of Theorem 1, the optimal allocation $q_i(\theta,s)$ is always increasing in k_i. Therefore, the sign of $\Delta_3 p_i$ depends only on $\frac{dk_i^*}{dk}$, which is characterized in Appendix Lemma 8; its sign depends on the comparison between the marginal investment cost and the capacity-sorting margin in EuJ. By contrast, in a contributing category, the effect of a capacity expansion may be ambiguous due to the additional revenue effect.

Figure 7 illustrates these two observations in a two-category environment. Category 1 has a lower average welfare weight and is therefore the contributing category, but it favors high types $(\tilde{\lambda}_1 = 0.9,\ a_1 = 1.6)$. Category 2 has a higher average welfare weight and is therefore the non-contributing category, but it favors low types $(\tilde{\lambda}_2 = 1.1,\ a_2 = -1.6)$. The qualitative pattern is driven by the ranking of average welfare weights. Note that swapping the direction of the preference tilt between the two categories, while holding average welfare weights fixed, produces the mirror image of the same pattern. The illustration assumes that k_i^* increases in both categories, that is, that the marginal gain in surplus in EuJ exceeds the marginal investment

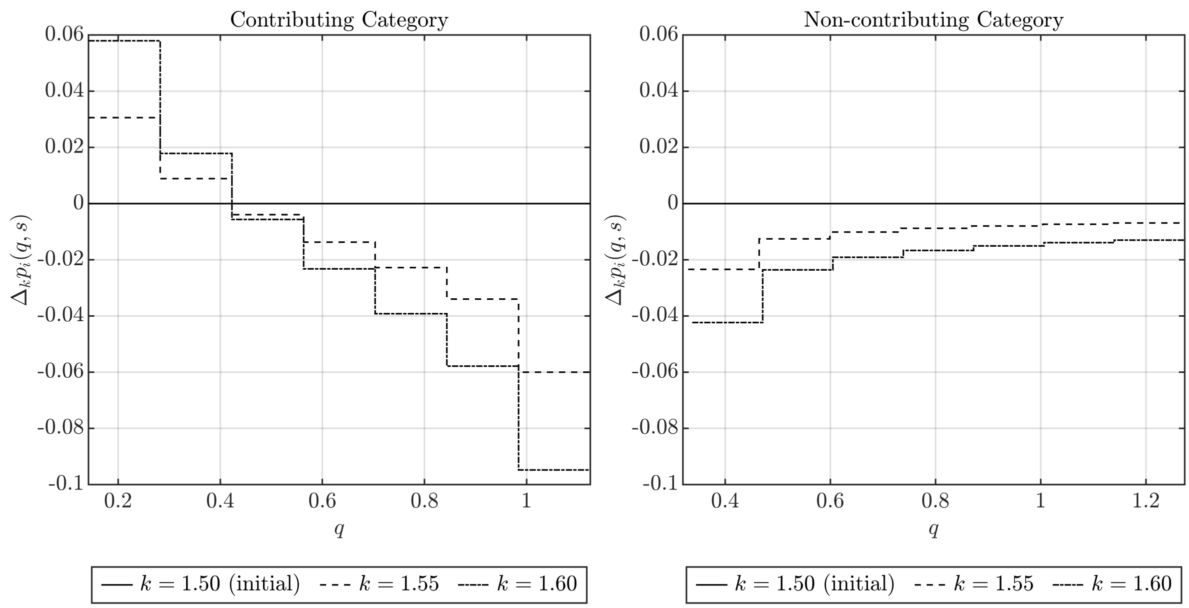

Figure 7: Difference in marginal price schedules by categories due to a capacity expansion.

cost. As expected, the marginal price is strictly decreasing for all types in the non-contributing category. For the contributing category, however, the change in the marginal price is not monotone in sign. In particular, lower types, despite being favored by the welfare weights, may face higher marginal prices after a capacity expansion. Appendix Figure 22 decomposes $\Delta_3 p_i$ into its two components and shows that this increase in marginal prices is driven by the capacity effect. In particular, as discussed after Theorem 1, the quantity allocated to lower types decreases because their contribution to surplus creation is negative; the formal expression is given in $\partial_k q_i$.

The incidence of individual surplus mirrors these price changes. Following Theorem 1, a capacity expansion increases surplus for all types in the non-contributing category. In the contributing category, by contrast, there exists a cutoff type such that lower types lose from the expansion, whereas higher types gain, despite the fact that welfare weights favor low types.

7 Conclusion

This article studies a peak-load pricing model in essential-goods sectors, such as electricity, transport, and other network industries, where capacity is limited in the short run and costly to expand. In this environment, a market designer must allocate scarce supply once demand is realized and decide ex ante how much capacity to build and how to finance its expansion. The problem is shaped by the coexistence of observable differences across consumer categories, unobservable heterogeneity within them, and redistributive objectives. In contrast to the standard peak-load pricing benchmark, this article characterizes who gains from capacity

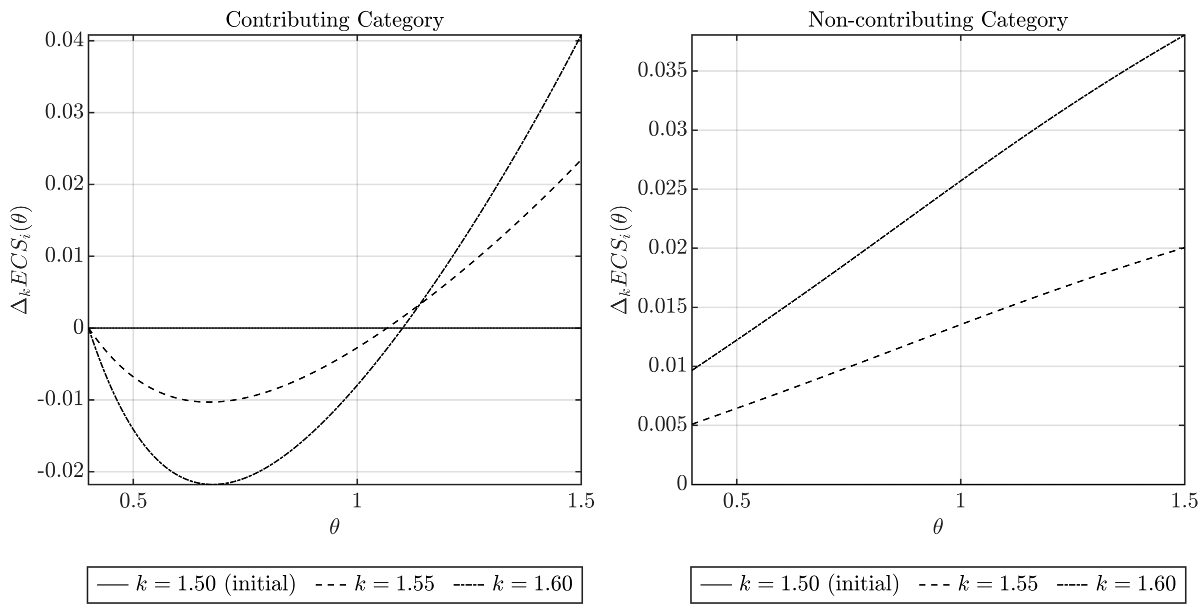

Figure 8: Difference in individual surplus by categories due to a capacity expansion.

expansions and how redistributive preferences reshape the allocation of consumption and surplus among consumers.

The answers to these questions are governed by two key elements: the change in the surplus generated by the mechanism and, therefore, available for redistribution through transfers; and the distribution of the social returns across consumers, which captures the redistributive and revenue motives at play in the environment. Because these objects need not align with the underlying redistributive priorities, it is not obvious either who gains from capacity expansion or which consumers benefit most from redistribution. To illustrate this tension, assume that a higher priority is given to consumers with low consumption. If their demand is more elastic, then relaxing scarcity can disproportionately increase the informational rents that must be left to higher types. These rents reduce the revenue that can be extracted from consumers with lower redistributive priority while having limited social value from the designer's perspective. As a result, the surplus available for redistribution may fall, reducing the surplus of consumers targeted by the redistribution.

These same elements also provide a unified framework to characterize the article's broader policy implications. In this setting, discrimination across observable categories, such as households with different incomes, is always welfare-improving, but it also generates distinct distributive effects. I therefore derive the optimal tagging rule for ranking categories by capacity allocation and budget contribution, and show that this ranking does not follow directly from redistributive priorities alone. In particular, the relevant tension depends on both the overall level of redistributive concern and the distribution of those preferences between low- and high-consumption consumers. Finally, redistributive preferences also distort both the optimal investment rule and the mechanism's tariff implementation compared to the utilitarian benchmark. As a result, spot prices fail to decentralize the optimum, while optimal nonlinear tariffs, including block tariffs, can take non-standard forms across categories and consumption levels, for the same reason that underlying redistributive priorities do not map directly into optimal allocations.

References

Ahlvik, L., Liski, M., and Mattila, M. (2024). Pigouvian Income Taxation. .

Akbarpour, M., Budish, E. B., Dworczak, P., and Kominers, S. D. (2023). An economic framework for vaccine prioritization. *Available at SSRN 3846931*.

Akbarpour, M., Dworczak, P., and Kominers, S. D. (2024). Redistributive allocation mechanisms. *Journal of Political Economy*, 132(6):1831–1875.

Akerlof, G. A. (1978). The economics of "tagging" as applied to the optimal income tax, welfare programs, and manpower planning. *American Economic Review*, 68(1):8–19.

Amores, A. F., Christl, M., Agostini, P. D., Poli, S. D., and Maier, S. (2025). Limiting prices or transferring money? an ex ante assessment of alternative measures to cope with the hike in energy prices. *Energy Economics*, 147.

Bierbrauer, F. (2024). Is a market-based approach to climate policy desirable? .

Boiteux, M. (1949). *La tarification des demandes en pointe: Application de la théorie de la vente au coût marginal*. place Henri-Bergson.

Borenstein, S. (2012). The redistributional impact of nonlinear electricity pricing. *American Economic Journal: Economic Policy*, 4:56–90.

Borenstein, S. (2025). Energy hogs and energy angels: What does residential electricity use really tell us about profligate consumption? *AEA Papers and Proceedings*, 115:357–62.

Burger, S. P., Knittel, C. R., Pérez-Arriaga, I. J., Schneider, I., and Vom Scheidt, F. (2020). The efficiency and distributional effects of alternative residential electricity rate designs. *Energy Journal*, 41(1):199–239.

Busby, J. W., Baker, K., Bazilian, M. D., Gilbert, A. Q., Grubert, E., Rai, V., Rhodes, J. D., Shidore, S., Smith, C. A., and Webber, M. E. (2021). Cascading risks: Understanding the 2021 winter blackout in texas.

Cahana, M., Fabra, N., Reguant, M., and Wang, J. (2022). The distributional impacts of real-time pricing. *CEPR Discussion Paper No. DP17200*.

Chao, H.-p. and Wilson, R. (1987). Priority service: Pricing, investment, and market organization. *The American Economic Review*, pages 899–916.

Chetty, R. and Saez, E. (2010). Optimal taxation and social insurance with endogenous private insurance. *American Economic Journal: Economic Policy*, 2(2):85–114.

Condorelli, D. (2013). Market and non-market mechanisms for the optimal allocation of scarce resources. *Games and Economic Behavior*, 82:582–591.

Cremer, H. and Gahvari, F. (2002). Nonlinear pricing, redistribution, and optimal tax policy. *Journal of Public Economic Theory*, 4(2):139–161.

Crew, M. A., Fernando, C. S., and Kleindorfer, P. R. (1995). The theory of peak-load pricing: A survey. *Journal of regulatory economics*, 8:215–248.

DallasFed (2021). Cost of texas' 2021 deep freeze justifies weatherization. https://www.dallasfed.org/research/economics/2021/0415. Dallas Fed Economics webpage, accessed April 3, 2026.

Dworczak, P., Kominers, S. D., and Akbarpour, M. (2021). Redistribution Through Markets. *Econometrica*, 89(4):1665–1698.

Ecola, L. and Light, T. (2009). Equity and Congestion Pricing A Review of the Evidence. Technical report, The RAND Corporation.

Eliasson, J. and Mattsson, L. G. (2006). Equity effects of congestion pricing. quantitative methodology and a case study for stockholm. *Transportation Research Part A: Policy and Practice*, 40:602–620.

EuropeanCommission (2025). Investment needs of european energy infrastructure to enable a decarbonised economy. Technical report, European Commission.

Feger, F., Pavanini, N., and Radulescu, D. (2022). Welfare and Redistribution in Residential Electricity Markets with Solar Power. *Review of Economic Studies*, 89(6):3267–3302.

Feldstein, M. S. (1972a). Distributional equity and the optimal structure of public prices. *The American Economic Review*, 62(1/2):32–36.

Feldstein, M. S. (1972b). Equity and efficiency in public sector pricing: The optimal two-part tariff*. *The Quarterly Journal of Economics*, 86(2):175–187.

Fetzer, T., Gazze, L., and Bishop, M. (2024). Distributional and climate implications of policy responses to energy price shocks. *Economic Policy*, 39:711–756.

Hall, J. D. (2021). Can tolling help everyone? estimating the aggregate and distributional consequences of congestion pricing. *Journal of the European Economic Association*, 19(1):441–474.

Hinchberger, A., Jacobsen, M., Knittel, C., Sallee, J., and van Benthem, A. (2024). The efficiency of dynamic electricity prices. Technical report, National Bureau of Economic Research.

Hjort, X.; Leflaive, M. (2020). Addressing the social consequences of tariffs for water supply and sanitation. *OECD Environmental Working Papers*, page 214.

IEA (2025). Renewables 2025 analysis and forecasts to 2030. Technical report, International Energy Agency.

Joskow, P. L. and Wolfram, C. D. (2012). Dynamic pricing of electricity. In *American Economic Review*, volume 102, pages 381–385.

Kang, Z. Y. (2021). Optimal Public Provision of Private Goods. *SSRN Electronic Journal*.

Kang, Z. Y. and Watt, M. (2024). Optimal Redistribution Through Subsidies.

Kleven, H. J. (2021). Sufficient Statistics Revisited. *Annual Review of Economics*, 13:515–538.

Levinson, A. and Silva, E. (2022). The Electric Gini: Income Redistribution through Energy Prices†. *American Economic Journal: Economic Policy*, 14(2):341–365.

Mäkimattila, M. (2025). Inequality-Aware Regulation. .

Manville, M., Pierce, G., and Graveline, B. (2022). Guardrails on priced lanes: Protecting equity while promoting efficiency. *Transportation Research Interdisciplinary Perspectives*, 15(August):100652.

Martimort, D., Pouyet, J., and Staropoli, C. (2020). Use and abuse of regulated prices in electricity markets: "How to regulate regulated prices?". *Journal of Economics and Management Strategy*, 29(3):605–634.

Maskin, E. and Riley, J. (1984). Monopoly with Incomplete Information. *RAND Journal of Economics*, 15(2):171–196.

Moore, S., Hill, E. M., Dyson, L., Tildesley, M. J., and Keeling, M. J. (2022). Retrospectively modeling the effects of increased global vaccine sharing on the covid-19 pandemic. *Nature Medicine*, 28:2416–2423.

Pai, M. and Strack, P. (2023). Taxing Externalities Without Hurting the Poor. *EC 2023 - Proceedings of the 24th ACM Conference on Economics and Computation*, page 1019.

Pathak, P. A., Sönmez, T., Ünver, M. U., and Yenmez, M. B. (2024). Fair allocation of vaccines, ventilators and antiviral treatments: Leaving no ethical value behind in healthcare rationing. *Management Science*, 70:3999–4036.

Randriamaro, M. T. and Cook, J. (2026). Income redistribution, water scarcity and water rates in the 225 largest us cities under the creative commons attribution 4.0 international license (cc by 4.0). *Journal of Water Economics*, 53:1–32.

Räsänen, M., Ruusunen, J., and Hämäläinen, R. P. (1997). Optimal tariff design under consumer self-selection. *Energy Economics*, 19(2):151–167.

Saez, E. (2001). Using Elasticities to Derive Optimal Income Tax Rates. *The Review of Economic Studies*, 68(1):205–229.

Schrank, D., Albert, L., Jha, K., and Eisele, B. (2025). 2025 urban mobility report. Technical report, Texas AM Transportation Institute, Tech. Rep.

Schweppe, F. C., Caramanis, M. C., Tabors, R. D., and Bohn, R. E. (1988). *Spot Pricing of Electricity*. Springer US, Boston, MA.

Smith, S. M. (2022). The effects of individualized water rates on use and equity. *Journal of Environmental Economics and Management*, 114.

Spence, M. (1977). Nonlinear prices and welfare. *Journal of Public Economics*, 8(1):1–18.

Spulber, D. F. (1992a). Capacity-contingent nonlinear pricing by regulated firms. *Journal of Regulatory Economics*, 4(4):299–319.

Spulber, D. F. (1992b). Optimal nonlinear pricing and contingent contracts. *International Economic Review*, pages 747–772.

Steiner, P. O. (1957). Peak Loads and Efficient Pricing. *The Quarterly Journal of Economics*, 71(4):585–610.

Vickrey, W. (1971). Responsive Pricing of Public Utility Services. *The Bell Journal of Economics and Management Science*, 2(1):337.

Vickrey, W. S. (1963). Pricing in urban and suburban transport. *The American Economic Review*, 53(2):452–465.

Vickrey, W. S. (1969). Congestion theory and transport investment. *The American Economic Review*, pages 251–260.

Wagner, N. (2018). The billion dollar question how much will it cost to decarbonise cities' transport systems? Technical report, International Transport Forum Discussion Papers, No. 2018/20, OECD Publishing.

Wang, D., Bjørnstad, O. N., Lei, T., Sun, Y., Huo, J., Hao, Q., Zeng, Z., Zhu, S., Hallegatte, S., Li, R., Guan, D., and Stenseth, N. C. (2023). Supply chains create global benefits from improved vaccine accessibility. *Nature Communications*, 14.

A Additional results

A.1 Single crossing of EuJ

This appendix establishes the single-crossing property of EuJ invoked in the main text, and provides the primitive foundations of Assumption Monot..

Definition 1. R_q *is* log-submodular (resp. log-supermodular) in (θ, k_i) *if* $\partial^2 \log R_q / \partial\theta\, \partial k_i \leq 0$ (*resp.* ≥ 0).

Lemma 6 (Single-crossing)**.** *Assume that, for each k_i, the map $\theta \mapsto R_\theta(\theta)$ is monotone, and that $R_q(\theta)$ has log-monotone differences in (θ, k_i). Then the map*

$$k_i \longmapsto \mathbb{E}^{w}_{(s,i)}[uJ \mid s]$$

is single-crossing in k_i. Moreover, if a sign change occurs, its direction is pinned down by the combination of monotonicities: it crosses from $+$ to $-$ when $R_\theta(\theta)$ is increasing and R_q is log-submodular, or when $R_\theta(\theta)$ is decreasing and R_q is log-supermodular; it crosses from $-$ to $+$ in the two remaining cases.

Proof. See Appendix C.3 $\square$

The sign of both ratios depends on the model primitives as follows. For $R'_q(\theta)$, only the curvature of u matters:

$$\text{sign}\Big(R'_q(\theta)\Big) = -\text{sign}\Big((u_q(\cdot))^2 - u(\cdot)u_{qq}(\cdot)\Big),$$

so $R'_q(\theta) < 0$ for linear and quadratic demand, while $R'_q(\theta) > 0$ for CES/isoelastic demand. For $R'_\theta(\theta)$, the answer depends on whether the virtual surplus $J_i(\theta)$ grows faster or slower than the IC term $\Lambda_i(\theta)$ as θ increases,

$$\text{sign}\big(R'_\theta(\theta)\big) = \text{sign}(J_i(\theta))\ \text{sign}\Bigg(\underbrace{\frac{J'_i(\theta)}{J_i(\theta)}}_{\text{revenue motive}} - \underbrace{\frac{\Lambda'_i(\theta)}{\Lambda_i(\theta)}}_{\text{redistributive motive}} \Bigg). \qquad (\text{sign}R_\theta)$$

For types contributing positively to the available surplus ($J_i(\theta) > 0$), $R_\theta(\theta)$ is increasing whenever the revenue motive grows proportionally faster than the redistributive motive. For types contributing negatively ($J_i(\theta) < 0$), the reverse holds: $R_\theta(\theta)$ is increasing whenever the revenue motive grows proportionally slower. When those assumptions fail, the same sufficient statistic remains valid, but its sign need not vary monotonically with k_i.[27]

On the IR-unconstrained region, where $\beta_i = 0$, the log-monotone-differences condition becomes a condition on primitives alone: since $\partial q_i / \partial k_i = w_i / \overline{w}_i$ is proportional across types, the sign of $\partial^2 \log R_q / \partial\theta\, \partial k_i$

[27] It is not possible to have a clear-cut sign of $R'_\theta(\theta)$. Indeed, the sign of $\Lambda'_i(\theta)$ depends on $g_i(\theta)$ and $\lambda_i(\theta)$: $\Lambda'_i(\theta) = -\Lambda_i(\theta)\frac{g'_i(\theta)}{g_i(\theta)} - \lambda_i(\theta)$. For instance, under a uniform distribution $\Lambda'_i(\theta) < 0$, while its sign may be ambiguous for other distributions. Hence, under a uniform distribution, the sign of $R_\theta(\theta)$ is ambiguous as soon as $J_i(\theta) < 0$.

reduces to the sign of $R_q''(q)$, so R_q is log-submodular whenever $R_q''(\cdot) \leq 0$ and log-supermodular whenever $R_q''(\cdot) \geq 0$. On the IR-constrained region, the response of the shadow cost β_i enters the derivative of $\partial q_i / \partial k_i$ and the condition is no longer implied by the curvature of U alone. Under Assumption CARA, however, $R_q = 1/\alpha$ is constant, so all cross-derivatives of $\log R_q$ vanish identically; the log-monotone-differences condition is therefore satisfied with equality regardless of the regime considered. Finally, note that unlike the log-monotone-differences condition on R_q, the monotonicity of $R_\theta(\theta)$ does not depend on k_i: from signR_θ, its sign is determined entirely by the type distribution and the welfare weights.

A.2 Three types results

A.2.1 Monotonicity of $\mathcal{G}_i(\theta)$

In this section, I show that in the two-type case, $\mathcal{G}_H > \mathcal{G}_L$ holds for any $\beta \geq 0$ and any welfare weights $\lambda_L, \lambda_H > 0$. This does not extend to three types: with three types $\theta_L < \theta_M < \theta_H$ and equal probabilities $p_I = 1/3$, the monotonicity of effective weights at any adjacent pair need not hold.

Two types. Using $J_H - J_L = (\theta_H - \theta_L)/p_L > 0$ and $\Gamma_H - \Gamma_L = \lambda_L(\theta_H - \theta_L) > 0$,

$$\mathcal{G}_H - \mathcal{G}_L = (\Gamma_H - \Gamma_L) + \beta(J_H - J_L) = (\lambda_L + \beta/p_L)(\theta_H - \theta_L) > 0.$$

Three types. Assume equal spacing $\theta_M - \theta_L = \theta_H - \theta_M = d$. A direct computation gives

$$\mathcal{G}_H - \mathcal{G}_M = d \cdot \frac{1}{3}\left[2(\lambda_L + \lambda_M) - \lambda_H + 6\beta\right],$$

$$\mathcal{G}_M - \mathcal{G}_L = d \cdot \frac{1}{3}\left[2(\lambda_L + \lambda_H) - \lambda_M + 6\beta\right].$$

At $\beta = 0$, each condition reduces to a simple primitive restriction:

$$\mathcal{G}_H > \mathcal{G}_M \iff \lambda_H < 2(\lambda_L + \lambda_M), \qquad \mathcal{G}_M > \mathcal{G}_L \iff \lambda_M < 2(\lambda_L + \lambda_H).$$

Either condition can fail depending on which welfare weight is disproportionately large relative to the other two.

A.2.2 Monotonicity of $R_\theta(\theta)$

In the two-type case with $\beta = 0$, R_θ is always increasing in θ: from $\Gamma_H = \theta_H \tilde{\lambda}$ and $\Gamma_L = J_L \tilde{\lambda} + \Lambda_L$, $R_\theta(\theta_H) - R_\theta(\theta_L) = \Lambda_L / (\tilde{\lambda}\Gamma_L) \geq 0$ under Assumption Int, with equality only when $\lambda_H = 0$. This monotonicity does not extend to three types.

Consider three types $\theta_L < \theta_M < \theta_H$ with probabilities p_L, p_M, p_H and welfare weights $\lambda_L, \lambda_M, \lambda_H$. Under $\beta = 0$, $R_\theta(\theta_H) = 1/\tilde{\lambda}$ regardless of the configuration, and $R_\theta(\theta_H) - R_\theta(\theta_M) = \Lambda_M/(\tilde{\lambda}\Gamma_M) \geq 0$, so the $M \to H$ step is always weakly increasing. The action lies at the $L \to M$ step. A direct computation gives

$$\text{sign}\big[R_\theta(\theta_M) - R_\theta(\theta_L)\big] = \text{sign}\big[J_M\Lambda_L - J_L\Lambda_M\big].$$

Substituting $J_L = \theta_L - \frac{p_M+p_H}{p_L}(\theta_M - \theta_L)$, $J_M = \theta_M - \frac{p_H}{p_M}(\theta_H - \theta_M)$, $\Lambda_L = \frac{(p_M\lambda_M+p_H\lambda_H)(\theta_M-\theta_L)}{p_L}$, and $\Lambda_M = \frac{p_H\lambda_H(\theta_H-\theta_M)}{p_M}$, the sign condition becomes

$$\text{sign}\big[\big(\theta_M p_M - p_H(\theta_H - \theta_M)\big)(p_M\lambda_M + p_H\lambda_H)(\theta_M - \theta_L) - \big(\theta_L p_L - (p_M + p_H)(\theta_M - \theta_L)\big)p_H\lambda_H(\theta_H - \theta_M)\big].$$

Specializing to equal probabilities $p_L = p_M = p_H = 1/3$, $\theta_H = 1$, and $\theta_M = \theta_L + d$, this reduces to

$$\text{sign}\big[d(2\theta_L + 2d - 1)(\lambda_M + \lambda_H) - (\theta_L - 2d)\lambda_H(1 - \theta_L - d)\big].$$

Non-monotonicity of R_θ requires this expression to be negative, which in turn requires $\theta_L > 2d$, i.e., types compressed toward the lower end. As d shrinks, the second term grows in absolute value while the first vanishes at order d, so narrow type spacing makes non-monotonicity easier to obtain, but also concentrates it in a narrow region of the parameter space.

A numerical example: take $p_L = p_M = p_H = 1/3$, $\theta_L = 0.8$, $\theta_M = 1.0$, $\theta_H = 1.5$, and welfare weights $\lambda_L = \lambda_M = 0.5$, $\lambda_H = 1$. Then $J_L = 0.4$, $J_M = 0.5$, $J_H = 1.5$, and $R_\theta(\theta_L) \approx 0.706$, $R_\theta(\theta_M) = 0.600$, $R_\theta(\theta_H) = 1.500$, so R_θ is non-monotone across the three types.

A.3 Implementation feasibility

This appendix collects the full characterization of the implementation region together with a numerical illustration.

Proposition 4 (Implementation feasibility)**.** *Fix category i and a budget requirement $I_i > 0$. Let k_i^+ be the capacity level at which the capacity constraint never binds.*

Assume that EuJ is single-crossing in k_i from $+$ to $-$ *on $[0, k_i^+]$. Then $\mathbb{E}\underline{CS}_i(k_i)$ is single-peaked on $[0, k_i^+]$ and the set of capacities for which the IR-unconstrained allocation is feasible (equivalently, $\mathbb{E}\underline{CS}_i(k_i) \geq 0$) has the following form:*

(i) If $\mathbb{E}\underline{CS}_i(k_i^+) \geq 0$, there exists a unique cutoff $\tilde{k}_i \in (0, k_i^+]$ such that the IR-unconstrained allocation is optimal for all $k_i \in [\tilde{k}_i, k_i^+]$.

(ii) If $\mathbb{E}\underline{CS}_i(k_i^+) < 0$*, the IR-unconstrained allocation is either never optimal on* $[0, k_i^+]$ *or optimal only on an intermediate interval* $[\tilde{k}_i^-, \tilde{k}_i^+] \subset (0, k_i^+)$*.*

If instead EuJ crosses from $-$ *to* $+$*, then* $\mathbb{E}\underline{CS}_i(k_i)$ *is single-dipped on* $[0, k_i^+]$ *and either* $\mathbb{E}\underline{CS}_i(k_i) < 0$ *for all* $k_i \in [0, k_i^+]$*, or there exists a unique* $\tilde{k}_i \in (0, k_i^+]$ *such that* $\mathbb{E}\underline{CS}_i(k_i) \geq 0$ *for all* $k_i \in [\tilde{k}_i, k_i^+]$*.*

Proof. See Appendix C.14 □

Proposition 4 identifies the geometry of the IR-unconstrained region. When EuJ switches from $+$ to $-$, the implementation region may arise only for intermediate values of k_i: low levels of capacity keep the financing constraint tight, while high levels of capacity may shift the gains from further expansion toward types that contribute relatively less to revenue. By contrast, when marginal expansions continue to increase the mechanism's revenue, the undistorted region expands as capacity increases.

Figure 9 illustrates Theorem 1 and Proposition 4. The left panel shows the boundary term $\mathbb{E}CS_i^\times$ as a function of k_i for two type supports, tracing both the implementation region and the sign of EuJ. The dashed curve has a lower $\underline{\theta}_i$, which increases the mass of consumers with $J_i < 0$ and shrinks the range of k_i for which EuJ is positive, shifting k^* to the left. The right panels illustrate the individual surplus effects at two capacity levels: at k_1, where EuJ is positive, all types gain from a capacity expansion; at k_2, where EuJ is negative, only types above the cutoff $\tilde{\theta}_i$ benefit.

A.4 Preference perturbations and effective social weights

This subsection provides primitive conditions on redistributive preference perturbations that ensure a well-behaved change in effective social weights across types. In particular, it characterizes when the induced change in effective weights has a constant sign, a property that is useful for deriving comparative statics in the main text.

Recall that $\Lambda_i(\theta)$ measures the social value of the information rent given to types above θ, evaluated at the welfare weights. I consider a perturbation of redistributive preferences $\lambda_i(\theta)$ and study how it affects the associated objects.

Lemma 7 (Preference shifts and effective social weights)**.** *Consider a weight function* $\lambda_i(\theta)$ *and its perturbation* $\lambda_i^\Delta(\theta)$*, with* $\Delta\lambda_i(\theta) := \lambda_i^\Delta(\theta) - \lambda_i(\theta)$ *and* $\Delta\mathcal{G}_i(\theta) := \mathcal{G}_i^\Delta(\theta) - \mathcal{G}_i(\theta)$*. Suppose* $\mathcal{G}_i(\theta), \mathcal{G}_i^\Delta(\theta) > 0$ *for all* $\theta \in [\underline{\theta}_i, \overline{\theta}_i]$*. Define:*

$$m_i(\theta) = \frac{\lambda_i^\Delta(\theta)}{\lambda_i(\theta)}, \qquad c_\Lambda := \frac{\tilde{\lambda}_i^\Delta}{\tilde{\lambda}_i}, \qquad c_\mathcal{G} := \frac{\tilde{\lambda}_i^\Delta + \beta_i^\Delta}{\tilde{\lambda}_i + \beta_i}.$$

If $m_i(\theta)$ is increasing (resp. decreasing) on $[\underline{\theta}_i, \overline{\theta}_i]$, then for $\theta \in [\underline{\theta}_i, \overline{\theta}_i]$,

$$\frac{\Lambda_i^{\Delta\prime}(\theta)}{\Lambda_i^{\Delta}(\theta)} \geq \text{ (resp. } \leq) \frac{\Lambda_i'(\theta)}{\Lambda_i(\theta)}.$$

Moreover, if $1 \leq c_{\mathcal{G}} \leq c_\Lambda$ (resp. $c_\Lambda \leq c_{\mathcal{G}} \leq 1$), then for all $\theta \in [\underline{\theta}_i, \overline{\theta}_i]$, $\Delta\mathcal{G}_i(\theta) \geq$ (resp. $\leq$)0. In particular, if either $\beta_i = \beta_i^\Delta = 0$ or $\beta_i^\Delta + \tilde{\lambda}_i^\Delta = \beta_i + \tilde{\lambda}_i$, then the condition reduces to $\tilde{\lambda}_i^\Delta \geq$ (resp. $\leq$)$\tilde{\lambda}_i$.

Proof. See Appendix C.9 □

The first part of the lemma shows that a proportional tilt of welfare weights toward higher types shifts $\Lambda_i^\Delta(\theta)$ upward relative to $\Lambda_i(\theta)$ across the type space. The second part connects the direction of a preference shift to the sign of $\Delta\mathcal{G}_i(\theta)$, under conditions that relate the average change in welfare weights to the corresponding change in effective weights. In particular, monotone mean-preserving rotations of $\lambda_i(\theta)$ generate a constant-sign change in $\Delta\mathcal{G}_i(\theta)$ across types.

A.5 Additional results on preference-induced reallocation

This appendix presents the technical results underlying the sorting effect induced by redistributive preference perturbations. The key identity is Lemma 3, which shows that, at fixed capacity,

$$\Delta q_i(\theta, s) = R_q(\theta)\left(R_{\mathcal{G}}(\theta) - \mathbb{E}_i^w[R_{\mathcal{G}}(\theta) \mid s]\right), \qquad s \in T_i.$$

Hence the quantity response depends on deviations of the proportional change in effective social weight from its capacity-weighted average.

Proposition 5 (Non-monotone reallocation under preference perturbations)**.** *Assume $R_{\mathcal{G}}(\theta)$ is not (a.e.) constant.*

(i) *Under the IR-unconstrained allocation, $\Delta q_i(\cdot, s)$ is not monotone on Θ_i for any perturbation. Under the IR-constrained allocation, assume that $\beta_i^\Delta + \tilde{\lambda}_i^\Delta = \beta_i + \tilde{\lambda}_i$, then the same conclusion holds whenever $|\Delta\tilde{\lambda}_i|$ is sufficiently small.*

(ii) *Assume a mean preserving perturbation such that $\Delta\tilde{\lambda}_i = 0$. Then $\Delta q_i(\cdot, s)$ is not monotone, the cutoff set $\{\theta : \Delta q_i(\theta, s) = 0\}$ contains at least two elements, and letting $\theta_- \leq \theta_+$ denote its minimum and maximum elements:*

(i) if $\Delta\mathcal{G}_i(\theta) \geq 0$ for all θ, then $\Delta q_i(\theta, s) \leq 0$ for $\theta \in [\underline{\theta}_i, \theta_-) \cup (\theta_+, \overline{\theta}_i]$;

(ii) if $\Delta\mathcal{G}_i(\theta) \leq 0$ for all θ, then $\Delta q_i(\theta, s) \geq 0$ for $\theta \in [\underline{\theta}_i, \theta_-) \cup (\theta_+, \overline{\theta}_i]$.

Proof. See Appendix C.15 □

Proposition 5 shows that the sorting effect is generically non-monotone. Even when preferences shift uniformly toward higher or lower types, the induced quantity response may not be ranked monotonically across the type space. Intuitively, the revenue motive, embedded in $J_i(\theta)$, and the redistributive motive, embedded in $\Lambda_i(\theta)$, interact through the endogenous scarcity adjustment. This forces reallocations toward some interior types at the expense of others, rather than producing a simple upward or downward tilt of the quantity schedule.

The second part sharpens the non-monotonicity result under mean-preserving perturbations. In that case, the boundary types necessarily move in the opposite direction from the interior region where the reallocation concentrates. Thus, even a preference shift that pointwise favors higher types need not raise the quantities received by the highest types; the IC structure instead pushes the adjustment toward intermediate types.

A.6 Marginal response of category allocations to aggregate capacity

The next lemma characterizes the local response of the optimal category allocation to a marginal change in aggregate capacity. The result is local: it holds on a neighborhood of an interior optimum in which the set of active and contributing categories is unchanged. Its main purpose is to show that cross-category heterogeneity enters only through the sufficient statistic $\mathbb{E}^{w}_{(s,i)}[uJ \mid s]$, which is the same object that governs the single-category comparative statics.

Lemma 8 (Marginal response of category allocations to aggregate capacity)**.** *Fix an interior optimum $\{k_i^*, I_i^*\}_{i\in N}$ and suppose that, on a neighborhood of the optimum, the partition $N = N^+ \cup N^0$ is unchanged. For each category $i \in N$, define*

$$\Omega_i := \mathbb{E}_s\left[\frac{1}{\overline{w}_i(s)}\mathbf{1}_{\{s\in T_i\}}\right] > 0, \qquad \hat{A}_i := \mathbb{E}^{w}_{(s,i)}[uJ \mid s]\,\mathbf{1}_{\{i\in N^+\}}.$$

Then the marginal response of category i is

$$\frac{dk_i^*}{dk} = \frac{\Omega_i^{-1}\big(\hat{A}_i - I'(k^*)\big)}{\sum_{j\in N^+}\mu_j\,\Omega_j^{-1}\,\mathbb{E}^{w}_{(s,j)}[uJ \mid s] \;-\; I'(k^*)\sum_{j\in N}\mu_j\,\Omega_j^{-1}}\,\mathcal{B} \;-\; \Omega_i^{-1}\tilde{\lambda}^* I''(k^*),$$

where

$$\mathcal{B} := 1 + \tilde{\lambda}^* I''(k^*)\sum_{j\in N}\mu_j\,\Omega_j^{-1} \;>\; 0$$

since $\tilde{\lambda}^* I''(k^*) > 0$. The formula takes two forms across the partition. For $i \in N^+$, the first-term numerator is $\Omega_i^{-1}\Big(\mathbb{E}^{w}_{(s,i)}[uJ \mid s] - I'(k^*)\Big)$. For $i \in N^0$, $\hat{A}_i = 0$ and the numerator reduces to $-\Omega_i^{-1}I'(k^*)$.

Proof. See Appendix C.16 □

The lemma shows that the response of k_i^* to aggregate expansion depends on category i only through the sufficient statistic $\mathbb{E}^{w}_{(s,i)}[uJ \mid s]$. This makes it possible to connect the outer comparative statics to the single-category ordering results.

B Additional Figure

Figure 9. $U(q,s) = (1+s)q - \frac{1}{4}q^2$, with $s \sim \mathcal{U}[0,2]$, $\theta \sim \mathcal{U}[0.10, 1.00]$, weights $\lambda_i(\theta) = 1 + \frac{\theta - 0.55}{0.90}$, and budget requirement $I_i = 0.45$. Under the uniform type distribution, $R_\theta(\theta)$ is increasing.

Figure 10. $U(q,s) = a(1+s)\frac{q^{1-\sigma}}{1-\sigma}$ with $\sigma = 0.6$ and $a = 0.0095$, with $s \sim \mathcal{U}[0,1]$ and $\theta \sim \mathcal{U}[1,2]$. The benchmark is $\lambda_i(\theta) = 1$, and the weights are $\lambda_i^A(\theta) = \frac{\exp(1.4(\theta-1.5))}{\int_1^2 \exp(1.4(t-1.5))dt}$ and $\lambda_i^B(\theta) = \frac{\exp(-1.4(\theta-1.5))}{\int_1^2 \exp(-1.4(t-1.5))dt}$. The allocation comparison is reported for $k_i = 0.90$, $s = 0.14$, and $I_i = 0$.

Figure 11. $\theta \sim \mathcal{U}[1,2]$, with utilitarian benchmark $\lambda(\theta) = 1$ and baseline effective social weight $\Gamma_0(\theta) = 2\theta - 1$. Welfare-weight perturbations take the form $\lambda_i(\theta) = 1 + \Delta(\theta)$, with $\Delta(\theta) = c + a\log(1 + 2(\theta - 1))$. The four cases are $A1 : (c,a) = (0.6765, -0.9546)$, $A2 : (c,a) = (0.1526, -0.2150)$, $B1 : (c,a) = (0.6765, -0.9546)$, and $B2 : (c,a) = (-0.4875, 0.6880)$. The figure reports the implied level effect $\Delta\Gamma(\theta)$ and proportional effect $R_\Gamma(\theta) = \Delta\Gamma(\theta)/\Gamma(\theta)$.

Figure 12. $U(q,s) = \frac{1+s}{5}\left(1 - e^{-5q}\right)$, with $s \sim \mathcal{U}[0,1]$ and $\theta \sim \mathcal{U}[0.3,2]$. There are $N = 50$ categories and aggregate capacity is fixed at $k = 30.00$. The initial welfare weight is $\lambda_1(\theta) = \exp\left(-0.60\left(\frac{\theta-0.3}{1.7} - \frac{1}{2}\right)\right)$. In both cases, the sequence $\{\lambda_i(\theta)\}_{i=1}^{50}$ is defined recursively by $\lambda_{i+1}(\theta) = \lambda_i(\theta)m_i(\theta)$, with $m_i(\theta) = \exp\left(c_i + a_i\left(\frac{\theta-0.3}{1.7} - \frac{1}{2}\right)\right)$, where c_i is chosen so that $\tilde{\lambda}_{i+1} = \tilde{\lambda}_i + 0.06$. Case A uses $a_i \in [0.01, 0.12]$, Case B uses $a_i \in [-0.12, -0.01]$, and $I_i = 0$.

Figure 13. $U(q,s) = \frac{1+s}{5}\left(1 - e^{-5q}\right)$, with $s \sim \mathcal{U}[0,1]$, $\theta \sim \mathcal{U}[0.2, 2.0]$, and $N = 2$. Aggregate capacity is evaluated over $k \in [0.50, 1.13]$, with investment cost $I(k) = 0.15\,k^2$. In the redistributive case, the initial welfare weight is $\lambda_1(\theta) = \exp\left(-0.10\left(\frac{\theta-0.2}{1.8} - \frac{1}{2}\right)\right)$, while $\lambda_2(\theta) \propto \lambda_1(\theta)\exp\left(0.02\left(\frac{\theta-0.2}{1.8} - \frac{1}{2}\right)\right)$, normalized so that $\tilde{\lambda}_2 = \tilde{\lambda}_1 + 0.06$. In the utilitarian benchmark, $\lambda_1(\theta) = \lambda_2(\theta) = 1$.

Figure 14. $U(q,s) = \frac{1+s}{5}\left(1 - e^{-5q}\right)$, with $s \sim \mathcal{U}[0,1]$, $\theta \sim \mathcal{U}[0.3, 2.0]$, and $N = 25$. For each successive pair $(i, i+1)$, aggregate capacity is chosen over $k \in [0.50, 1.50]$, with investment cost $I(k) = 0.15\,k$. The initial welfare weight is $\lambda_1(\theta) = 1$, and the sequence $\{\lambda_i(\theta)\}_{i=1}^{25}$ is defined recursively by $\lambda_{i+1}(\theta) = \lambda_i(\theta)m_i(\theta)$, where $m_i(\theta) = \exp\left(c_i + a_i\left(\frac{\theta-0.3}{1.7} - \frac{1}{2}\right)\right)$ and c_i is chosen so that $\tilde{\lambda}_i = 1$ for all categories. Case A uses $a_i = 0.005 + 0.015\,t_i^2$, while Case B uses $a_i = -0.005 - 0.015\,t_i^2$, with t_i evenly spaced on $[0,1]$.

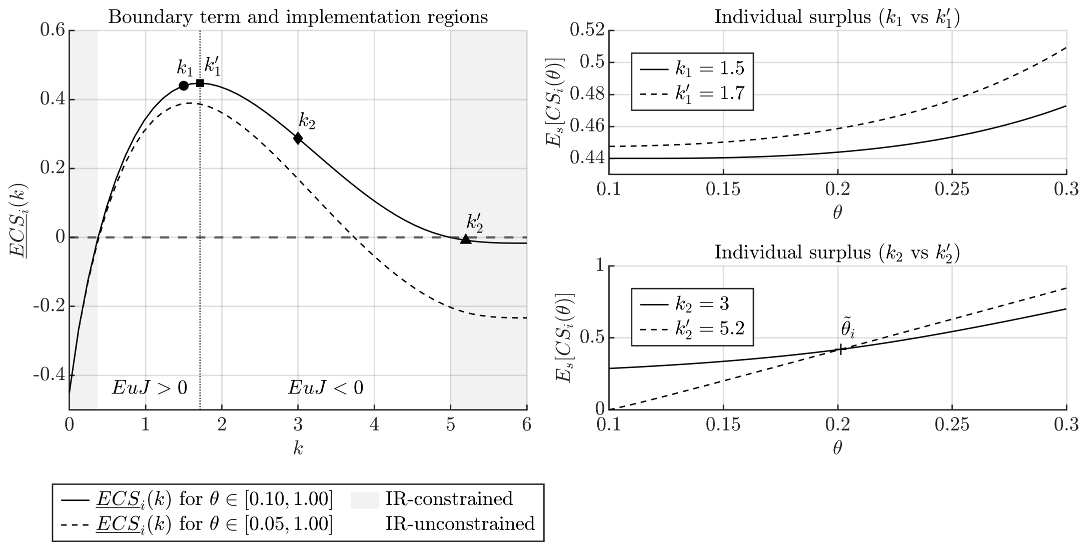

Figure 9: Individual surplus incidence of a capacity expansion. Left panel: surplus available to the designer for redistribution as a function of the category's allocated capacity. Right panel: individual surplus across types under the optimal allocation, for different levels of capacity.

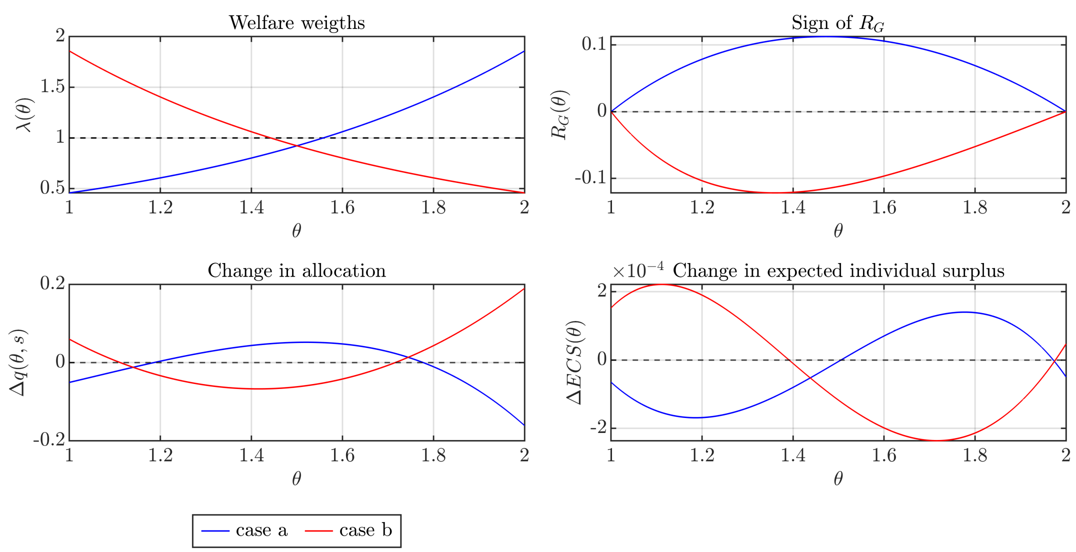

Figure 10: Allocation and surplus incidence of redistributive preference perturbations relative to the utilitarian benchmark. First panel: welfare weights. Second panel: proportional change in effective social weights. Third panel: difference in the optimal allocation for each type relative to the benchmark. Fourth panel: difference in expected individual surplus.

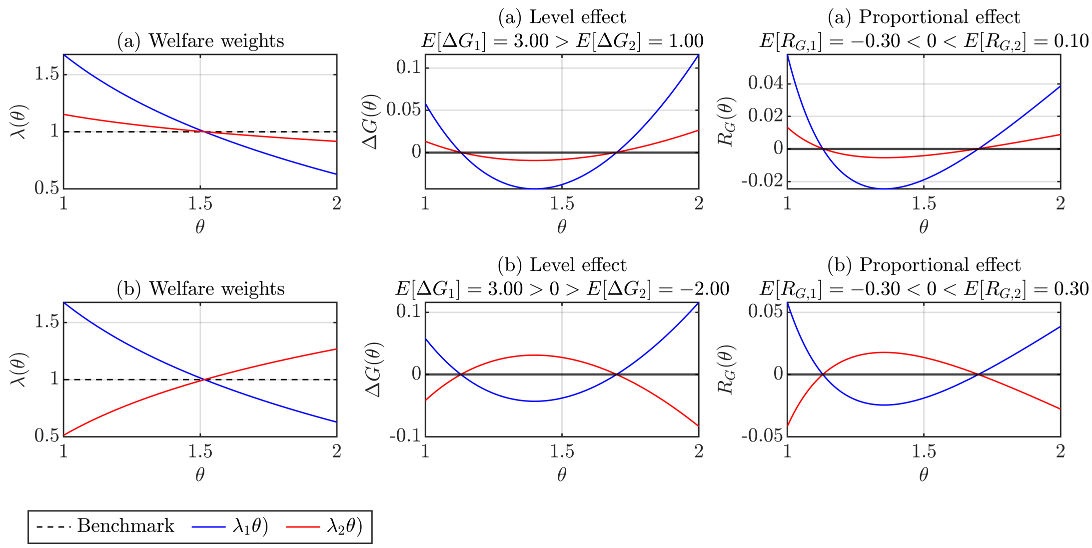

Figure 11: Incidence of welfare weights perturbations relative to the utilitarian benchmark in terms of absolute and relative social weights. First column: welfare weights. Second column: absolute difference in effective social weights aggregated at the category level. Third column: proportional difference in effective social weights aggregated at the category level.

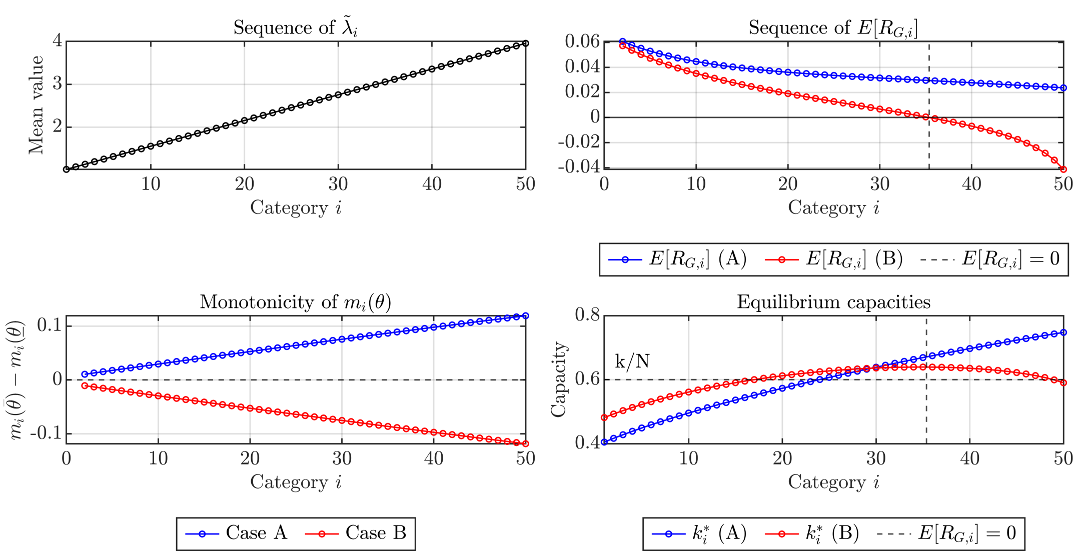

Figure 12: Optimal allocation of aggregate capacity ($k = 30$) across 50 categories. First panel: average social weights. Second panel: proportional difference in effective social weights aggregated at the category level. Third panel: indicator of which types are preferred (positive: higher types, negative: lower types). Fourth panel: Equilibrium capacity allocated to each category.

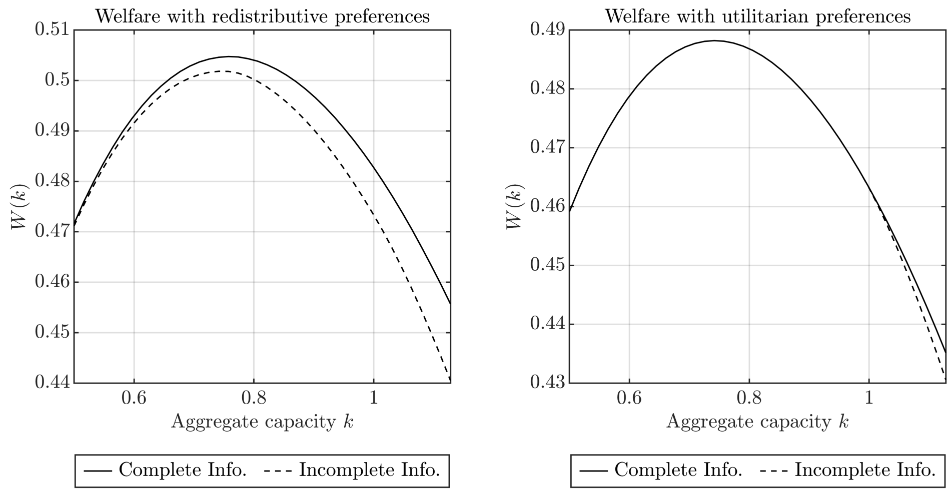

Figure 13: Optimal welfare as a function of aggregate capacity k. Left panel: welfare with and without private information under redistributive preferences. Right panel: utilitarian preferences.

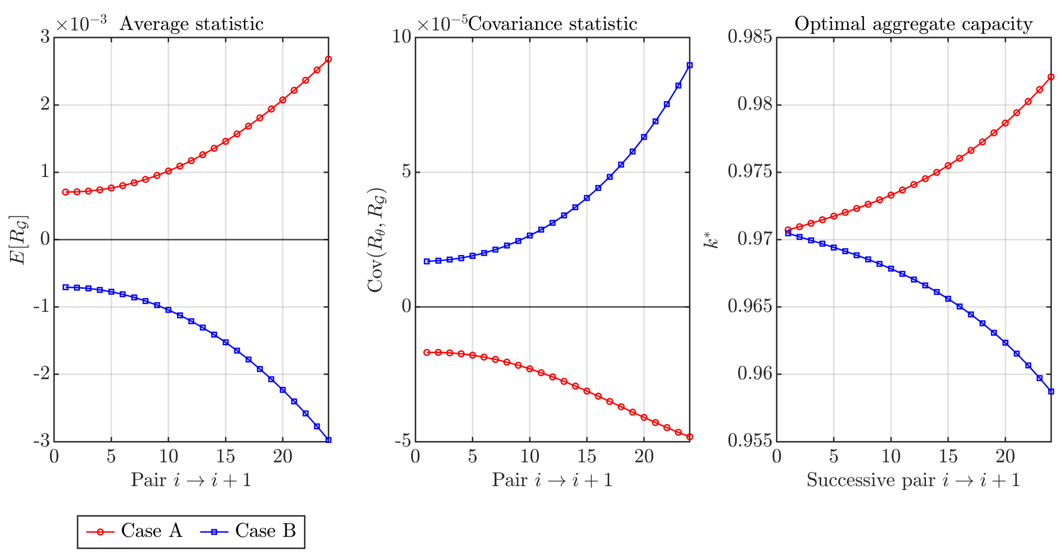

Figure 14: Optimal investment level under different redistributive preferences. First panel: first term in Proposition 3. Second panel: second term in Proposition 3. Third panel: optimal investment level.

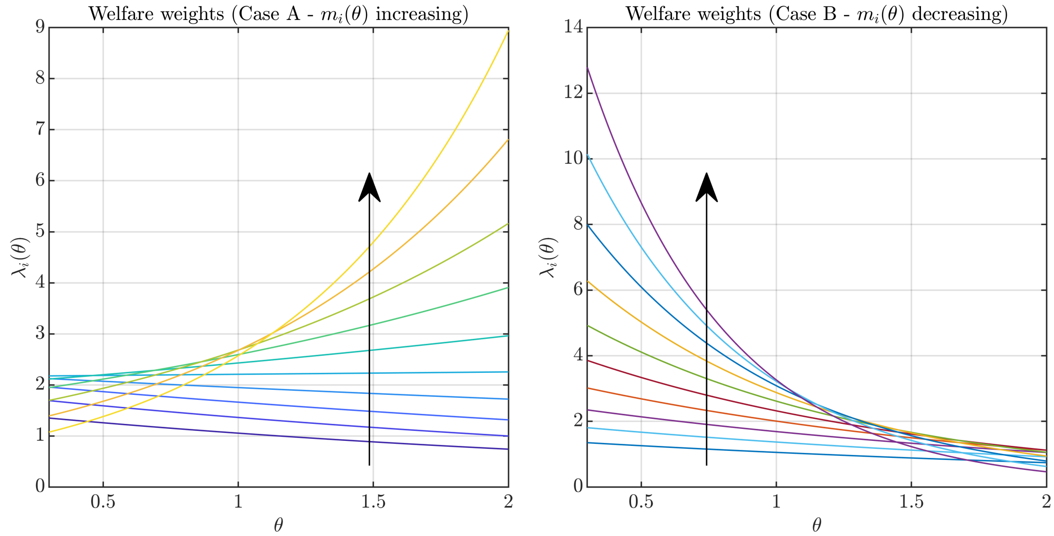

Figure 15: Welfare weights for each category in the example of Figure 12.

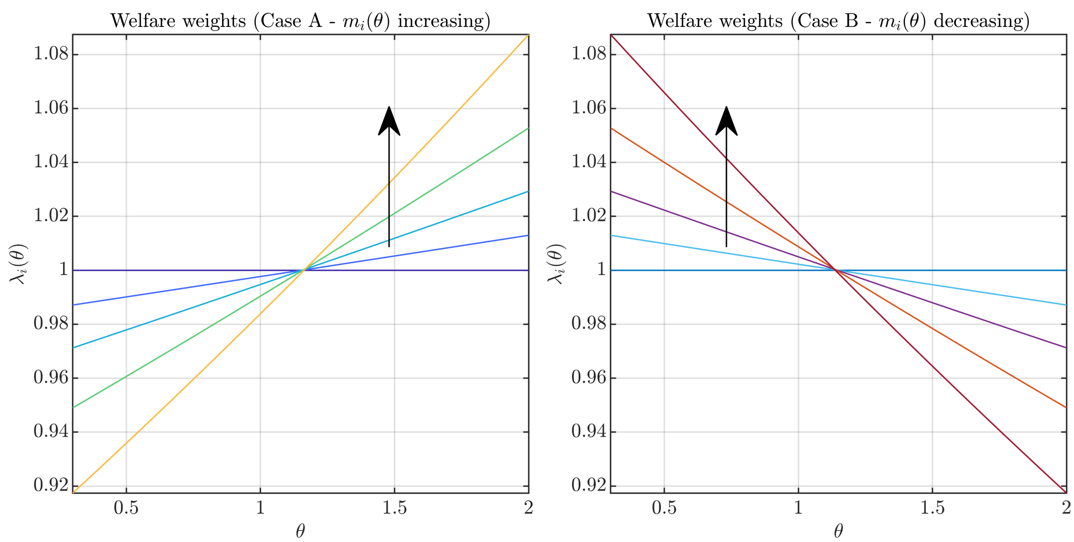

Figure 16: Welfare weights for each category in the example of Figure 14.

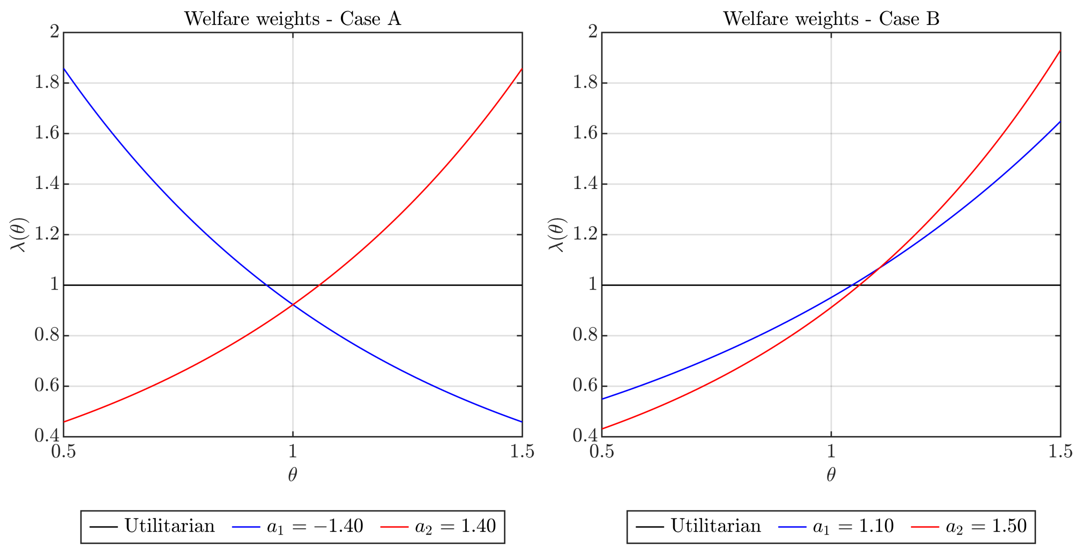

Figure 17: Welfare-weight profiles for each category.

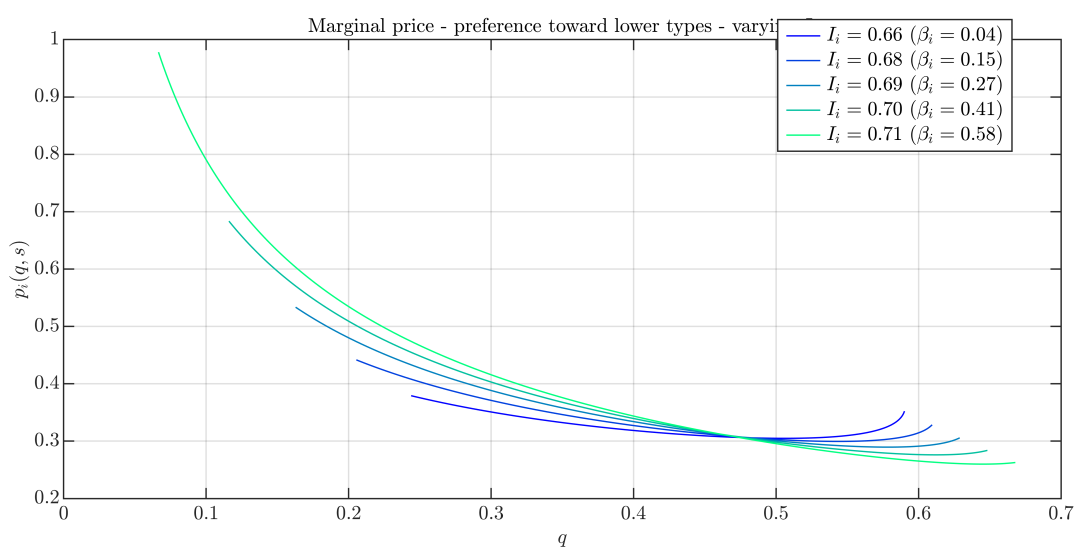

Figure 18: Marginal price schedules for different revenue requirements I_i.

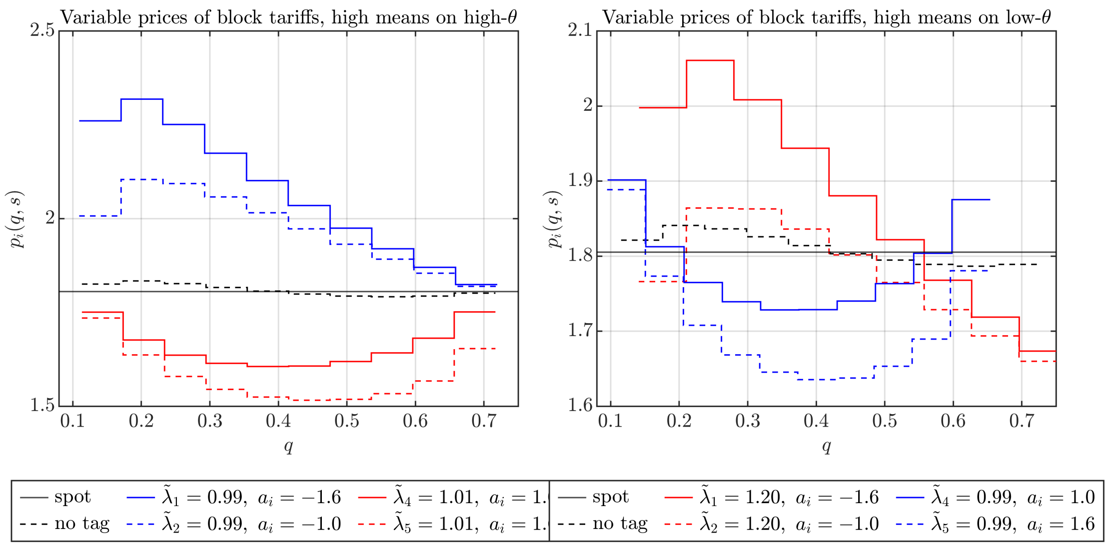

Figure 19: Marginal price schedules for the two other cases, with comparison to the common spot-price and no-tagging benchmarks.

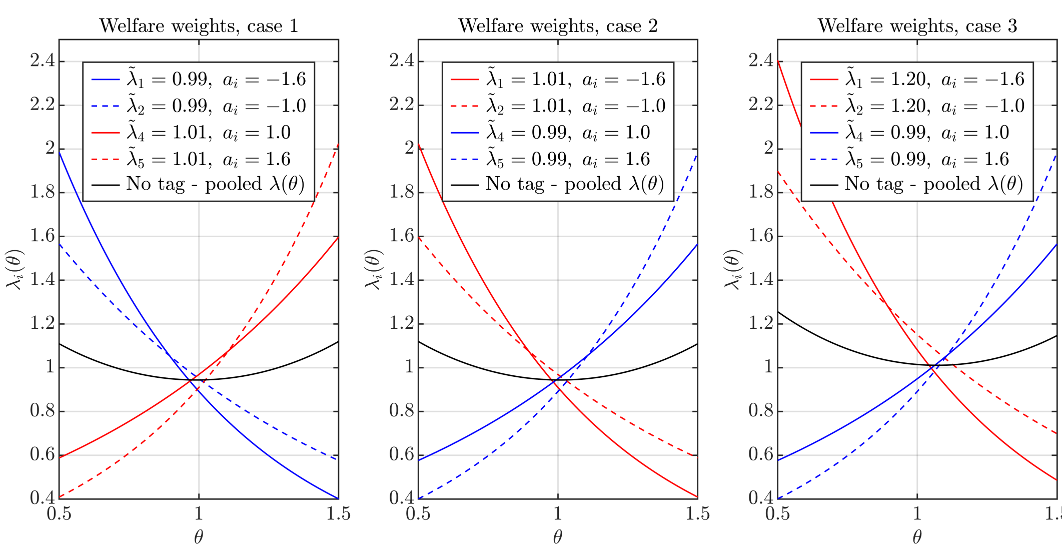

Figure 20: Category-specific welfare-weight profiles used in the three across-category illustrations, together with the pooled profile under no tagging.

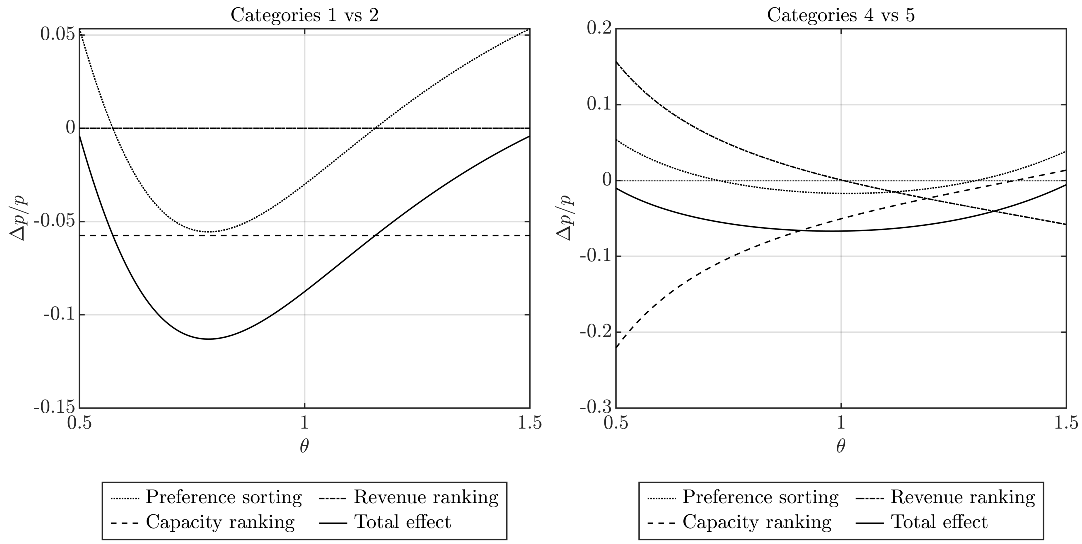

Figure 21: Decomposition of the relative marginal price difference across types, for two pairs of adjacent categories.

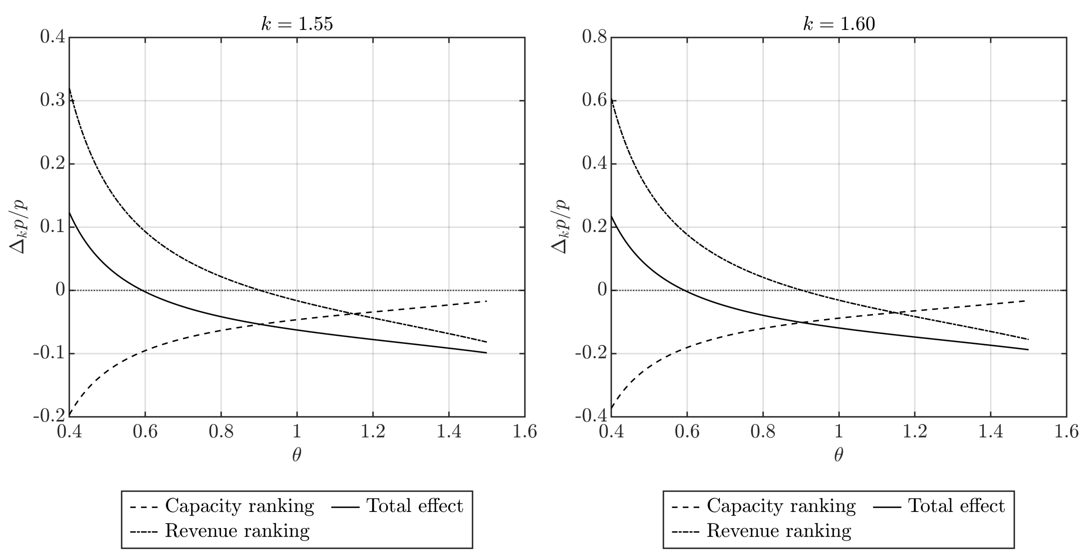

Figure 22: Decomposition of the relative marginal price difference after capacity expansion.

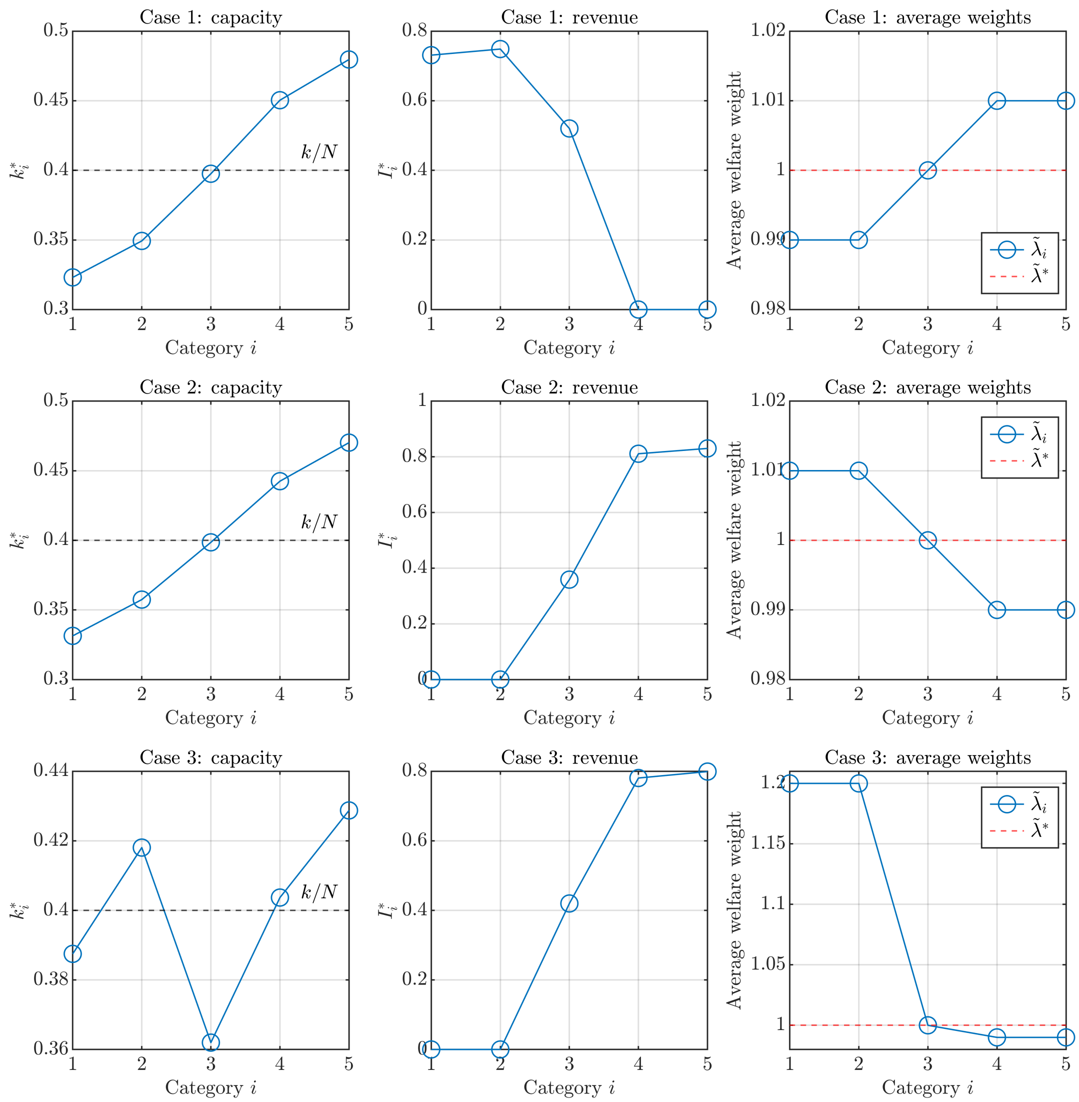

Figure 23: Optimal category-level allocations of capacity, revenue requirements, and average welfare weights in the three across-category illustrations.

C Proofs

C.1 Virtual surplus representation - Proof of Lemma 1

Proof. For each category i, define

$$\Phi_i(\theta, s) := \theta U(q_i(\theta, s), s) - \underline{\theta}_i U(q_i(\underline{\theta}_i, s), s) - \int_{\underline{\theta}_i}^{\theta} U\big(q_i(\tilde{\theta}, s), s\big)\, d\tilde{\theta},$$

and let $\mathbb{E}\Phi_i(\theta) := \mathbb{E}_s[\Phi_i(\theta, s)]$.

Start with the following individual transfer for each consumer of type θ in category i:

$$t_i(\theta, s) = \tau_i + \mathbb{E}\Phi_i(\theta) - \sum_{j \neq i} \frac{\mu_j}{1 - \mu_j}\, \mathbb{E}_j[\mathbb{E}\Phi_j(\theta)] + \Big(I(k) - I_i + \sum_{j \neq i} \frac{\mu_j}{1 - \mu_j} I_j\Big),$$

where

$$\tau_i = \underline{t}_i + \sum_{j \neq i} \frac{\mu_j}{1 - \mu_j}\, \mathbb{E}_j[\mathbb{E}\Phi_j(\theta)] - \Big(I(k) - I_i + \sum_{j \neq i} \frac{\mu_j}{1 - \mu_j} I_j\Big),$$

where $\underline{t}_i$ is the transfer of the lowest type $\underline{\theta}_i$. Taking expectations over θ, and noting that only $\mathbb{E}\Phi_i(\theta)$ depends on θ, yields

$$\mathbb{E}_i[t_i(\theta, s)] = \tau_i + \mathbb{E}_i[\mathbb{E}\Phi_i(\theta)] - \sum_{j \neq i} \frac{\mu_j}{1 - \mu_j}\, \mathbb{E}_j[\mathbb{E}\Phi_j(\theta)] + \Big(I(k) - I_i + \sum_{j \neq i} \frac{\mu_j}{1 - \mu_j} I_j\Big).$$

Multiplying by the size μ_i of category i, summing over all categories, and using $\sum_{i=1}^n \mu_i = 1$, we obtain

$$\begin{aligned}\sum_{i=1}^n \mu_i \mathbb{E}_i[t_i(\theta, s)] = \sum_{i=1}^n \mu_i \tau_i + \sum_{i=1}^n \mu_i\, \mathbb{E}_i[\mathbb{E}\Phi_i(\theta)] - \sum_{i=1}^n \mu_i \sum_{j \neq i} \frac{\mu_j}{1 - \mu_j}\, \mathbb{E}_j[\mathbb{E}\Phi_j(\theta)] \\ + I(k) - \sum_{i=1}^n \mu_i I_i + \sum_{i=1}^n \mu_i \sum_{j \neq i} \frac{\mu_j}{1 - \mu_j} I_j.\end{aligned}$$

The second and third terms cancel, and the fifth and sixth terms also cancel, so that $\sum_{i=1}^n \mu_i \mathbb{E}_i[t_i(\theta, s)] = \sum_{i=1}^n \mu_i \tau_i + I(k)$. Therefore, an allocation satisfies the budget constraint if and only if $\sum_{i=1}^n \mu_i \tau_i \geq 0$.

Next, substitute the expression for τ_i:

$$\begin{aligned}\sum_{i=1}^n \mu_i \tau_i &= \sum_{i=1}^n \mu_i \left(\underline{t}_i + \sum_{j \neq i} \frac{\mu_j}{1 - \mu_j}\, \mathbb{E}_j[\mathbb{E}\Phi_j(\theta)] - \Big(I(k) - I_i + \sum_{j \neq i} \frac{\mu_j}{1 - \mu_j} I_j\Big)\right) \\ &= \sum_{i=1}^n \mu_i \left(\underline{t}_i + \mathbb{E}_i[\mathbb{E}\Phi_i(\theta)]\right) - I(k),\end{aligned}$$

By Fubini's theorem and the usual hazard-rate argument, the expectation of $\mathbb{E}\Phi_i(\theta)$ can be written in terms of the virtual utility $U(q_i(\theta,s),s)J_i(\theta) := U(q_i(\theta,s),s)\,(\theta - \gamma_i(\theta)\,)$, as

$$\mathbb{E}_i[\mathbb{E}\Phi_i(\theta)] = \mathbb{E}_{(s,i)}\left[U(q_i(\theta,s),s)J_i(\theta) - \underline{\theta}_i U(q_i(\underline{\theta}_i,s),s)\right].$$

Hence

$$\sum_{i=1}^{n}\mu_i\tau_i = \sum_{i=1}^{n}\mu_i\mathbb{E}_{(s,i)}\left[\underline{t}_i + U(q_i(\theta,s),s)J_i(\theta) - \underline{\theta}_i U(q_i(\underline{\theta}_i,s),s)\right] - I(k).$$

Finally, replacing τ_i in $t_i(\theta,s)$ yields

$$t_i(\theta,s) = \underline{t}_i + \mathbb{E}\Phi_i(\theta) = \underline{t}_i + \mathbb{E}_s\left[\theta U(q_i(\theta,s),s) - \underline{\theta}_i U(q_i(\underline{\theta}_i,s),s) - \int_{\underline{\theta}_i}^{\theta} U\big(q_i(\tilde{\theta},s),s\big)\; d\tilde{\theta}\right].$$

Thus

$$\begin{aligned}\mathbb{E}_s\big[\theta U(q_i(\theta,s),s) - t_i(\theta,s)\big] = \mathbb{E}_s\big[\underline{\theta}_i U(q_i(\underline{\theta}_i,s),s)\big] - \underline{t}_i \\ + \int_{\underline{\theta}_i}^{\theta}\mathbb{E}_s\big[U(q_i(\tilde{\theta},s),s)\big]\,d\tilde{\theta}.\end{aligned}$$

Defining $\mathbb{E}\underline{CS}_i := \mathbb{E}_s\big[\underline{\theta}_i U(q_i(\underline{\theta}_i,s),s)\big] - \underline{t}_i$, we get

$$\mathbb{E}_s[\theta U(q_i(\theta,s),s) - t_i(\theta,s)] = \mathbb{E}\underline{CS}_i + \int_{\underline{\theta}_i}^{\theta}\mathbb{E}_s\big[U(q_i(\tilde{\theta},s),s)\big]\,d\tilde{\theta},$$

which is the usual envelope representation and therefore satisfies the IC condition. Moreover, substituting $\underline{\theta}_i U(q_i(\underline{\theta}_i,s),s) = \mathbb{E}\underline{CS}_i + \underline{t}_i$ into the expression for $\sum_i \mu_i\tau_i$ gives the claimed transformed budget constraint

$$\sum_{i=1}^{n}\mu_i\tau_i = \sum_{i=1}^{n}\mu_i\mathbb{E}_{(s,i)}\big[U(q_i(\theta,s),s)J_i(\theta) - \mathbb{E}\underline{CS}_i\big] - I(k).$$

Multiplying the utility of a consumer with type θ from category i and taking the expectation over the whole category yields

$$\mathbb{E}_{(s,i)}[\lambda_i(\theta)\;(\;\theta U(q_i(\theta,s),s) - t_i(\theta,s)\;)\;] = \tilde{\lambda}_i\;\mathbb{E}\underline{CS}_i + \mathbb{E}_i[\lambda_i(\theta)\int_{\underline{\theta}_i}^{\theta}\mathbb{E}_s\big[U(q_i(\tilde{\theta},s),s)\big]\,d\tilde{\theta}],$$

We apply Fubini's theorem to the second term:

$$\mathbb{E}_i\left[\lambda_i(\theta)\int_{\underline{\theta}_i}^{\theta}\mathbb{E}_s[U(q_i(\tilde{\theta},s),s)]\,d\tilde{\theta}\right] = \int_{\underline{\theta}_i}^{\bar{\theta}_i}\lambda_i(\theta)g_i(\theta)\left(\int_{\underline{\theta}_i}^{\theta}\mathbb{E}_s[U(q_i(\tilde{\theta},s),s)]\,d\tilde{\theta}\right)d\theta$$
$$= \int_{\underline{\theta}_i}^{\bar{\theta}_i}\mathbb{E}_s[U(q_i(\tilde{\theta},s),s)]\left(\int_{\tilde{\theta}}^{\bar{\theta}_i}\lambda_i(\theta)g_i(\theta)\,d\theta\right)d\tilde{\theta}$$
$$= \int_{\underline{\theta}_i}^{\bar{\theta}_i}\mathbb{E}_s[U(q_i(\tilde{\theta},s),s)]\gamma_i(\tilde{\theta})g_i(\tilde{\theta})\mathbb{E}_i[\,\lambda_i(\theta)\mid\theta\;\geq\;\tilde{\theta}\,]d\tilde{\theta}$$
$$= \mathbb{E}_{(s,i)}[\Lambda_i(\theta)\,U(q_i(\theta,s),s)]\,.$$

This yields the desired expression:

$$\tilde{\lambda}_i\ \mathbb{E}\underline{CS}_i + \mathbb{E}_{(s,i)}\left[\ \Lambda_i(\theta)\ U(q_i(\theta,s),s)\ \right]$$

□

C.2 Off-peak and on-peak allocation - Proof of Lemma 2

Because the revenue-feasibility constraint involves the virtual-surplus term $J_i(\theta)$, which may change sign across types, the convexity of the single-category feasible set is not automatic. We therefore maintain the following condition.

Assumption 7 (Convexity of the single-category feasible set)**.** *For every category i and every $I_i \geq 0$, the set $\left\{q_i(\theta,s) : \mathbb{E}_{(s,i)}[J_i(\theta)U(q_i(\theta,s),s)] \geq I_i\right\}$ is convex.*

Under Assumption 7, the single-category problem is a concave program, so the KKT conditions used in the proof of Lemma 2 are sufficient.

Proof. Consider the single-category problem $\max_{q_i(\theta,s)}\mathbb{E}_{(s,i)}[\Gamma_i(\theta)\,U(q_i(\theta,s),s)]$ subject to

$$\mathbb{E}_{(s,i)}[U(q_i(\theta,s),s)\,J_i(\theta)] - I_i \geq 0, \qquad \mathbb{E}_i[q_i(\theta,s)] \leq k_i \quad \forall s.$$

Let $\beta_i \geq 0$ and $\varepsilon_i(s) \geq 0$ denote the multipliers on IR-R$_i$ and the capacity constraint K$_i$, respectively. The associated Lagrangian is

$$\mathcal{L} = \mathbb{E}_{(s,i)}[(\Gamma_i(\theta) + \beta_i J_i(\theta))\,U(q_i(\theta,s),s)] - \mathbb{E}_s[\varepsilon_i(s)\,(\mathbb{E}_i[q_i(\theta,s)] - k_i)] - \beta_i I_i.$$

Hence the pointwise first-order condition is

$$u(q_i(\theta,s),s)\,\mathcal{G}_i(\theta) - \varepsilon_i(s) = 0, \qquad \text{where } \mathcal{G}_i(\theta) := \Gamma_i(\theta) + \beta_i J_i(\theta).$$

We first characterize the states in which the capacity constraint is slack. If $\mathbb{E}_i[q_i(\theta,s)] < k_i$, complementary slackness implies $\varepsilon_i(s) = 0$. Since Assumption Int gives $\mathcal{G}_i(\theta) > 0$ for all θ, the first-order condition reduces to $u(q_i(\theta,s),s) = 0$. By strict concavity of $U(\cdot,s)$, this equation has a unique solution, denoted $q_i^0(s)$. In particular, the off-peak allocation is independent of θ, so it automatically satisfies monotonicity.

Moreover, implicit differentiation of $u(q_i^0(s),s) = 0$ yields

$$\frac{\partial q_i^0(s)}{\partial s} = -\frac{u_s\big(q_i^0(s),s\big)}{u_q\big(q_i^0(s),s\big)} \geq 0,$$

since $u_s \geq 0$ and $u_q < 0$. Therefore the unconstrained aggregate demand $Q_i^0(s) := \mathbb{E}_i[q_i^0(s)]$ is nondecreasing in s. It follows that the set of states in which the capacity constraint is slack is an interval of the form $S_i = [\underline{s}, s_i)$, while the set of binding states is $T_i = [s_i, \bar{s}]$, where s_i is defined by $\mathbb{E}_i[q_i(\theta,s_i)] = k_i$, with the obvious boundary cases if one of the two sets is empty.

Now consider $s \in T_i$. For such states, the allocation solves the state-by-state problem with binding aggregate quantity constraint, and the first-order condition is exactly $u(q_i(\theta,s),s)\,\mathcal{G}_i(\theta) = \varepsilon_i(s)$. Because $u_q < 0$ and $\mathcal{G}_i(\theta) > 0$, the objective is strictly concave in $q_i(\theta,s)$ pointwise, so this first-order condition characterizes a unique allocation for each (θ,s).

Finally, Assumption Int also imposes that $\mathcal{G}_i(\theta)$ is nondecreasing. Differentiating the first-order condition with respect to θ gives

$$\frac{\partial q_i(\theta,s)}{\partial \theta} = -\frac{u(q_i(\theta,s),s)}{u_q(q_i(\theta,s),s)}\,\frac{\mathcal{G}_i'(\theta)}{\mathcal{G}_i(\theta)} \geq 0,$$

so the monotonicity constraint is satisfied.

□

C.3 Single-crossing - Proof of Lemma 6

Proof. For every binding state $s \in T_i(k_i)$, define the probability measure $dP_{k_i,s}(\theta) := \frac{R_q(\theta)\,g_i(\theta)}{\mathbb{E}_i[R_q(\theta)]}\,d\theta$. Since $R_\theta(\theta) = \frac{J_i(\theta)}{\mathcal{G}_i(\theta)}$, we can write

$$\mathbb{E}_i^w[uJ \mid s] = \frac{\mathbb{E}_i\big[R_q(\theta)\,R_\theta(\theta)\big]}{\mathbb{E}_i\big[R_q(\theta)\big]} = \int_{\Theta_i} R_\theta(\theta)\,dP_{k_i,s}(\theta).$$

Differentiating $\mathbb{E}_i^w[uJ \mid s]$ with respect to k_i gives $\frac{\partial \mathbb{E}_i^w[uJ|s]}{\partial k_i} = M_i(k_i, s) + D_i(k_i, s)$, where

$$M_i(k_i, s) := \frac{\partial}{\partial t} \int_{\Theta_i} R_\theta(\theta)\, dP_{t,s}(\theta) \Big|_{t=k_i}$$

is the pure measure-shift term, and $D_i(k_i, s) := \int_{\Theta_i} \frac{\partial R_\theta(\theta)}{\partial k_i}\, dP_{k_i,s}(\theta)$ is the direct effect through the k_i-dependence of $R_\theta(\theta)$.

By the log-monotone-differences assumption on R_q, the family $\{P_{k_i,s}\}_{k_i}$ is ordered by monotone likelihood ratio for each binding state s. Hence, since $\theta \mapsto R_\theta(\theta)$ is monotone for each k_i, $M_i(k_i, s) \geq 0$ in the two cases

$$R_\theta(\cdot, k_i) \text{ increasing and } R_q \text{ log-supermodular}, \qquad R_\theta(\cdot, k_i) \text{ decreasing and } R_q \text{ log-submodular},$$

while $M_i(k_i, s) \leq 0$ in the two opposite cases.

Next, since

$$R_\theta(\theta) = \frac{J_i(\theta)}{\mathcal{G}_i(\theta)}, \qquad \mathcal{G}_i(\theta) = \Gamma_i(\theta) + J_i(\theta)\beta_i(k_i),$$

we have

$$\frac{\partial R_\theta(\theta)}{\partial k_i} = -\frac{\partial \beta_i(k_i)}{\partial k_i} \frac{J_i(\theta)^2}{\mathcal{G}_i(\theta)^2}.$$

Therefore $D_i(k_i, s) = -\frac{\partial \beta_i(k_i)}{\partial k_i}\, \mathcal{Q}_i(k_i, s)$, where

$$\mathcal{Q}_i(k_i, s) := \int_{\Theta_i} \frac{J_i(\theta)^2}{\mathcal{G}_i(\theta)^2}\, dP_{k_i,s}(\theta) > 0.$$

Thus

$$\frac{\partial \mathbb{E}_i^w[uJ \mid s]}{\partial k_i} = M_i(k_i, s) - \frac{\partial \beta_i(k_i)}{\partial k_i}\, \mathcal{Q}_i(k_i, s).$$

We now relate $\frac{\partial \beta_i(k_i)}{\partial k_i}$ to $\mathbb{E}_{(s,i)}^w[uJ \mid s]$. On the IR-constrained region:

$$\mathbb{E}\underline{CS}_i(k_i, \beta) = \mathbb{E}_{(s,i)}[U(q_i(\theta, s; k_i, \beta), s)\, J_i(\theta)] - I_i.$$

By construction, $\mathbb{E}\underline{CS}_i(k_i, \beta_i(k_i)) = 0$.

Holding β fixed, the standard comparative-statics identity gives, for every binding state s, $\frac{\partial q_i(\theta,s;k_i,\beta)}{\partial k_i} = \frac{w_i(\theta,s)}{\overline{w}_i(s)}$. Hence

$$\frac{\partial \mathbb{E}\underline{CS}_i(k_i, \beta)}{\partial k_i} = \mathbb{E}_s\Big[\mathbf{1}_{\{s \in T_i(k_i)\}}\, \mathbb{E}_i^w[uJ \mid s]\Big] = \mathbb{E}_{(s,i)}^w[uJ \mid s], \tag{1}$$

where the moving-threshold term vanishes because $u(q_i(\theta, s_i(k_i)), s_i(k_i)) = 0 \qquad \text{for every } \theta.$

Holding k_i fixed, differentiating the pointwise FOC $u(q_i(\theta,s),s)\,\mathcal{G}_i(\theta) = \varepsilon_i(s)$ with respect to β, using $\frac{\partial \mathcal{G}_i(\theta)}{\partial \beta} = J_i(\theta)$ and imposing $\mathbb{E}_i\left[\frac{\partial q_i(\theta,s)}{\partial \beta}\right] = 0$, yields

$$\frac{\partial \varepsilon_i(s)}{\partial \beta} = \mathbb{E}_i^w[uJ \mid s], \tag{2}$$

$$\frac{\partial q_i(\theta,s)}{\partial \beta} = w_i(\theta,s)\,\left(u(q_i(\theta,s),s)\,J_i(\theta) - \mathbb{E}_i^w[uJ \mid s]\right). \tag{3}$$

Therefore

$$\frac{\partial \mathbb{E}\underline{CS}_i(k_i,\beta)}{\partial \beta} = \mathbb{E}_s\left[\mathbf{1}_{\{s\in T_i(k_i)\}}\,\overline{w}_i(s)\,\mathrm{Var}_i^w(uJ \mid s)\right] =: \mathbb{E}_{(s,i)}^w[\mathrm{Var}_i(uJ)], \tag{4}$$

where $\mathrm{Var}_i^w(uJ \mid s)$ is the associated weighted variance such that $\mathrm{Var}_i^w(X \mid s) := \mathbb{E}_i^w\left[\left(X - \mathbb{E}_i^w[X \mid s]\right)^2 \,\middle|\, s\right]$ and $\mathbb{E}_{(s,i)}^w[\mathrm{Var}_i(uJ)] > 0$ away from knife-edge cases. Differentiating $\mathbb{E}\underline{CS}_i(k_i,\beta_i(k_i)) = 0$ with respect to k_i and using 1 and 4 gives

$$\frac{\partial \beta_i(k_i)}{\partial k_i} = -\frac{\mathbb{E}_{(s,i)}^w[uJ \mid s]}{\mathbb{E}_{(s,i)}^w[\mathrm{Var}_i(uJ)]}. \tag{5}$$

Finally, differentiate and using Leibniz' rule,

$$\frac{\partial \mathbb{E}_{(s,i)}^w[uJ \mid s]}{\partial k_i} = -\mathbb{E}_i^w[uJ \mid s](k_i,s_i(k_i))\,f(s_i(k_i))\,\frac{\partial s_i(k_i)}{\partial k_i} + \int_{s_i(k_i)}^{\bar{s}} \frac{\partial \mathbb{E}_i^w[uJ \mid s]}{\partial k_i}\,f(s)\,ds.$$

At the threshold state $s_i(k_i)$ we have $u(q_i(\theta,s_i(k_i)),s_i(k_i)) = 0 \qquad$ for every θ, so $\mathbb{E}_i^w[uJ \mid s](k_i,s_i(k_i)) = 0$. Hence the boundary term vanishes and

$$\frac{\partial \mathbb{E}_{(s,i)}^w[uJ \mid s]}{\partial k_i} = \overline{M}_i(k_i) - \frac{\partial \beta_i(k_i)}{\partial k_i}\,\overline{\mathcal{Q}}_i(k_i),$$

where

$$\overline{M}_i(k_i) := \int_{s_i(k_i)}^{\bar{s}} M_i(k_i,s)\,f(s)\,ds, \qquad \overline{\mathcal{Q}}_i(k_i) := \int_{s_i(k_i)}^{\bar{s}} \mathcal{Q}_i(k_i,s)\,f(s)\,ds > 0.$$

Now let k_i^0 satisfy $\mathbb{E}_{(s,i)}^w[uJ \mid s](k_i^0) = 0$. Then $\frac{\partial \beta_i(k_i^0))}{\partial k_i} = 0$, so $\frac{\partial \mathbb{E}_{(s,i)}^w[uJ|s]}{\partial k_i}(k_i^0) = \overline{M}_i(k_i^0)$. Since each $M_i(k_i^0,s)$ has the same sign across binding states, $\overline{M}_i(k_i^0)$ has that same sign. Therefore every sign-changing zero of $\mathbb{E}_{(s,i)}^w[uJ \mid s]$ is crossed in the same direction: $\frac{\partial \mathbb{E}_{(s,i)}^w[uJ|s]}{\partial k_i}(k_i^0) > 0$ when

$$R_\theta(\cdot,k_i) \text{ increasing and } R_q \text{ log-supermodular}, \qquad R_\theta(\cdot,k_i) \text{ decreasing and } R_q \text{ log-submodular},$$

so any sign change is from $-$ to $+$; and $\frac{\partial \mathbb{E}_{(s,i)}^w[uJ|s]}{\partial k_i}(k_i^0) < 0$ in the two opposite cases, so any sign change is from $+$ to $-$.

Thus the map $k_i \longmapsto \mathbb{E}_{(s,i)}^w[uJ \mid s]$ is single-crossing in k_i, with crossing direction pinned down by the combination of monotonicities stated in the lemma. $\square$

C.4 Redistributive effects of capacity expansion - Proof of Theorem 1

Proof. Define $\mathcal{W}_i(\theta) := \partial \mathbb{E}CS_i(\theta)/\partial k_i$. From the envelope representation $\mathbb{E}CS_i^{\times}$,

$$\mathcal{W}_i(\theta) = \mathcal{W}_i(\underline{\theta}_i) + \int_{\underline{\theta}_i}^{\theta} \mathcal{W}_i'(\tilde{\theta})\, d\tilde{\theta}, \qquad \mathcal{W}_i'(\theta) = \mathbb{E}_s\left[u(q_i(\theta,s),s)\,\frac{\partial q_i(\theta,s)}{\partial k_i}\mathbf{1}_{\{s\in T_i\}}\right].$$

For each binding state $s \in T_i$, differentiating the pointwise FOC FOC_q with respect to k_i and imposing the binding capacity constraint K gives

$$\frac{\partial q_i(\theta,s)}{\partial k_i} = \frac{w_i(\theta,s)}{\overline{w}_i(s)} - w_i(\theta,s)\,\mathbb{E}_i^w[uJ \mid s]\,\frac{\partial \beta_i(k_i)}{\partial k_i} + w_i(\theta,s)\,u(q_i(\theta,s),s)\,J_i(\theta)\,\frac{\partial \beta_i(k_i)}{\partial k_i},$$

where w_i, $\overline{w}_i$, and $\mathbb{E}_i^w[uJ \mid s]$ are as defined in the proof of Lemma 6.

IR-unconstrained allocation. Since $\beta_i = 0$, we have $\frac{\partial \beta_i(k_i)}{\partial k_i} = 0$ and $\mathcal{G}_i(\theta) \equiv \Gamma_i(\theta)$. The expression for $\partial q_i/\partial k_i$ reduces to $w_i(\theta,s)/\overline{w}_i(s)$, so

$$\mathcal{W}_i'(\theta) = \mathbb{E}_s\left[\frac{u(q_i(\theta,s),s)\,w_i(\theta,s)}{\overline{w}_i(s)}\mathbf{1}_{\{s\in T_i\}}\right] = \mathbb{E}_s\left[\frac{R_q(\theta)}{\Gamma_i(\theta)\,\overline{w}_i(s)}\mathbf{1}_{\{s\in T_i\}}\right] > 0,$$

since $R_q > 0$, $\Gamma_i(\theta) > 0$, and $\overline{w}_i(s) > 0$ by Assumption Int. Hence $\mathcal{W}_i$ is strictly increasing on $[\underline{\theta}_i, \overline{\theta}_i]$. By 1, $\mathcal{W}_i(\underline{\theta}_i) = \partial \mathbb{E}\underline{CS}_i/\partial k_i$ has the sign of EuJ. If EuJ is positive, $\mathcal{W}_i(\underline{\theta}_i) > 0$ and strict monotonicity implies every type gains. If EuJ is negative, $\mathcal{W}_i(\underline{\theta}_i) < 0$ and, by strict monotonicity, $\mathcal{W}_i$ crosses zero at most once; if and only if $\mathcal{W}_i(\overline{\theta}_i) > 0$, a unique cutoff $\tilde{\theta}_i$ exists and types below lose while types above gain.

IR-constrained allocation.

Step 1: initial condition. Since $\mathbb{E}\underline{CS}_i = 0$ holds identically along the IR-constrained region because $\beta_i(k_i)$ is chosen to maintain the binding IR , its total derivative with respect to k_i is zero. By the envelope representation $\partial_k CS$, this total derivative equals $\mathcal{W}_i(\underline{\theta}_i)$, so $\mathcal{W}_i(\underline{\theta}_i) = 0$.

Step 2: shape of $\mathcal{W}_i$. Under Assumption Multi, $u(q,s) = \kappa(q)\zeta(s)$ so the ratio $R_q(\theta) := -\kappa(q_i(\theta,s))/\kappa_q(q_i(\theta,s))$ is s-independent. Using the first-order condition FOC_q, we have $u(q_i(\theta,s),s)\,w_i(\theta,s) = R_q(\theta)/\mathcal{G}_i(\theta)$ and $u(q_i(\theta,s),s)\,J_i(\theta) = \varepsilon_i(s)\,R_\theta(\theta)$. Substituting into $\mathcal{W}_i'(\theta)$ yields

$$\mathcal{W}_i'(\theta) = \frac{R_q(\theta)}{\mathcal{G}_i(\theta)}\left(\Omega_i - \frac{\partial \beta_i(k_i)}{\partial k_i}\mathbb{E}_{(s,i)}^w[uJ \mid s] + \frac{\partial \beta_i(k_i)}{\partial k_i}\overline{\varepsilon}_i R_\theta(\theta)\right),$$

where $\Omega_i := \mathbb{E}_s\left[\overline{w}_i(s)^{-1}\mathbf{1}_{\{s\in T_i(k_i)\}}\right] > 0$ and $\overline{\varepsilon}_i := \mathbb{E}_s[\varepsilon_i(s)\mathbf{1}_{\{s\in T_i\}}] > 0$. Since 5 gives $\frac{\partial \beta_i(k_i)}{\partial k_i} = -\frac{\mathbb{E}_{(s,i)}^w[uJ|s]}{\mathbb{E}_{(s,i)}^w[\text{Var}_i(uJ)]}$, the term $-\frac{\partial \beta_i(k_i)}{\partial k_i}\mathbb{E}_{(s,i)}^w[uJ \mid s] = \frac{\mathbb{E}_{(s,i)}^w[uJ|s]^2}{\mathbb{E}_{(s,i)}^w[\text{Var}_i(uJ)]} \geq 0$, so the constant part of the bracket is strictly positive. Since $R_q(\theta)/\mathcal{G}_i(\theta) > 0$, the monotonicity of $\mathcal{W}_i'(\theta)$ is governed entirely by the term $\frac{\partial \beta_i(k_i)}{\partial k_i}\overline{\varepsilon}_i R_\theta(\theta)$: if $R_\theta(\theta)$ is monotone, $\mathcal{W}_i'(\theta)$ is monotone in θ, so $\mathcal{W}_i$ is either concave or convex. The sign of $\frac{\partial \beta_i(k_i)}{\partial k_i}$ is opposite to the sign of EuJ, so if $R_\theta(\theta)$ is increasing, $\mathcal{W}_i$ is concave when EuJ is positive and convex when EuJ is negative; if $R_\theta(\theta)$ is decreasing, these conclusions are reversed.

Step 3: existence and direction of the second cutoff. Since $\mathcal{W}_i(\underline{\theta}_i) = 0$, a concave $\mathcal{W}_i$ crosses zero at most once in the interior. If $\mathcal{W}_i'(\underline{\theta}_i) > 0$, the function initially rises from zero, eventually returns toward zero, and crosses it at a unique interior cutoff $\tilde{\theta}_i \in (\underline{\theta}_i, \overline{\theta}_i)$; types below $\tilde{\theta}_i$ gain and types above lose in case (i), and the reverse in case (ii). If $\mathcal{W}_i'(\underline{\theta}_i) \leq 0$, no interior cutoff exists and either all types lose or all types gain. The convex case is symmetric, with concavity and convexity interchanged, and all surplus inequalities reversed. If $R_\theta(\theta)$ is decreasing, concavity and convexity are swapped and all inequalities reverse throughout. □

C.5 Reallocation under preference perturbations - Proof of Lemma 3

Proof. Fix $s \in T_i$ and a local perturbation $\Delta\lambda_i(\theta)$ holding k_i fixed.

Using the weights $w_i(\theta, s)$ and $\mathbb{E}_i^w[\cdot \mid s]$ from the preliminary step of Lemma 6, note that $R_q(\theta, s) = \varepsilon_i(s)\, w_i(\theta, s)$, which follows from

$$R_q(\theta, s) = -\frac{u(q_i, s)}{u_q(q_i, s)} \qquad \text{and} \qquad \varepsilon_i(s) = u(q_i, s)\mathcal{G}_i(\theta).$$

Linearizing the pointwise first-order condition $u(q_i(\theta, s), s)\mathcal{G}_i(\theta) = \varepsilon_i(s)$ gives

$$u_q(q_i(\theta, s), s)\, \mathcal{G}_i(\theta)\Delta q_i(\theta, s) + u(q_i(\theta, s), s)\, \Delta\mathcal{G}_i(\theta) - \Delta\varepsilon_i(s) = 0. \tag{6}$$

Solving 6 for $\Delta q_i(\theta, s)$ and substituting the definitions of R_q, $R_\mathcal{G}(\theta) = \Delta\mathcal{G}_i(\theta)/\mathcal{G}_i(\theta)$, and $\varepsilon_i(s)$ yields

$$\Delta q_i(\theta, s) = R_q(\theta, s)\left(R_\mathcal{G}(\theta) - \frac{\Delta\varepsilon_i(s)}{\varepsilon_i(s)}\right). \tag{7}$$

Since the capacity constraint binds state-by-state and k_i is held fixed, $\mathbb{E}_i[\Delta q_i(\theta, s)] = 0$ for each $s \in T_i$. Substituting 7, using $R_q(\theta, s) = \varepsilon_i(s) w_i(\theta, s)$, dividing by $\varepsilon_i(s)\overline{w}_i(s) > 0$, and rearranging gives

$$\frac{\Delta\varepsilon_i(s)}{\varepsilon_i(s)} = \mathbb{E}_i^w[R_\mathcal{G}(\theta) \mid s]\,. \tag{8}$$

Substituting 8 into 7 gives

$$\Delta q_i(\theta, s) = R_q(\theta, s)\left(R_\mathcal{G}(\theta) - \mathbb{E}_i^w[R_\mathcal{G}(\theta) \mid s]\right), \tag{9}$$

which is the desired identity. □

C.6 Surplus and preference perturbations - Proof of Proposition 1

Proof. The surplus change follows from the IC representation $\mathbb{E}CS_i^{\times}$:

$$\Delta\mathbb{E}CS_i(\theta) = \Delta\mathbb{E}\underline{CS}_i + \int_{\underline{\theta}_i}^{\theta} \mathbb{E}_s\left[u(q_i(\tilde{\theta},s),s)\,\Delta q_i(\tilde{\theta},s)\,\mathbf{1}_{\{s\in T_i\}}\right] d\tilde{\theta}.$$

Substituting 9 and using the first-order condition $u(q_i,s) = \varepsilon_i(s)/\mathcal{G}_i(\theta)$ together with $R_q(\theta) = 1/\alpha$ under Assumption CARA, the derivative with respect to θ is

$$\frac{\partial\Delta\mathbb{E}CS_i(\theta)}{\partial\theta} = \frac{1}{\alpha}\mathbb{E}_s\left[\frac{\varepsilon_i(s)}{\mathcal{G}_i(\theta)}\left(R_{\mathcal{G}}(\theta) - \mathbb{E}_i[R_{\mathcal{G}}(\theta)]\right)\mathbf{1}_{\{s\in T_i\}}\right] = \frac{\bar{\varepsilon}_i}{\alpha\,\mathcal{G}_i(\theta)}\left(R_{\mathcal{G}}(\theta) - \mathbb{E}_i[R_{\mathcal{G}}(\theta)]\right),$$

where $\bar{\varepsilon}_i := \mathbb{E}_s[\varepsilon_i(s)\mathbf{1}_{\{s\in T_i\}}] > 0$. Since $\bar{\varepsilon}_i/(\alpha\,\mathcal{G}_i(\theta)) > 0$, the sign of the derivative is determined solely by $R_{\mathcal{G}}(\theta) - \mathbb{E}_i[R_{\mathcal{G}}(\theta)]$. Under unimodality of $R_{\mathcal{G}}(\theta)$, this expression crosses zero at most twice, so $\Delta\mathbb{E}CS_i(\theta)$ has at most two turning points and at most three zeros in $[\underline{\theta}_i, \overline{\theta}_i]$.

Boundary conditions.

IR-constrained allocation. The IR constraint $\mathbb{E}\underline{CS}_i = 0$ binds at both the baseline and the perturbed allocation, so $\Delta\mathbb{E}\underline{CS}_i = 0$ and therefore $\Delta\mathbb{E}CS_i(\underline{\theta}_i) = 0$.

IR-unconstrained allocation. Linearizing $\mathbb{E}CS_i^{\times}$ with respect to the preference perturbation and substituting 9:

$$\Delta\mathbb{E}\underline{CS}_i = \mathbb{E}_{(s,i)}\left[u(q_i,s)\,\Delta q_i(\theta,s)\,J_i(\theta)\,\mathbf{1}_{\{s\in T_i\}}\right] = \mathbb{E}_{(s,i)}\left[\varepsilon_i(s)\,R_\theta(\theta)\,R_q(\theta)\,(R_{\mathcal{G}}(\theta) - \mathbb{E}_i[R_{\mathcal{G}}(\theta)])\,\mathbf{1}_{\{s\in T_i\}}\right].$$

Under Assumption CARA, $\mathbb{E}_i^w[R_{\mathcal{G}}(\theta) \mid s] = \mathbb{E}_i[R_{\mathcal{G}}(\theta)]$ is s-independent, so the expression separates and yields

$$\Delta\mathbb{E}\underline{CS}_i = \frac{\bar{\varepsilon}_i}{\alpha}\,\mathrm{Cov}_i(R_\theta(\theta), R_{\mathcal{G}}(\theta)).$$

Cutoff structure. Since $R_{\mathcal{G}}(\theta)$ is unimodal, the derivative $\frac{\partial\Delta\mathbb{E}CS_i(\theta)}{\partial\theta} \propto R_{\mathcal{G}}(\theta) - \mathbb{E}_i[R_{\mathcal{G}}(\theta)]$ changes sign at most twice, so $\Delta\mathbb{E}CS_i(\theta)$ has at most two turning points and hence at most three interior zeros. Under $\Delta\mathcal{G}_i(\theta) \geq 0$, $R_{\mathcal{G}}(\theta)$ peaks above its mean in the interior so the derivative pattern is $-, +, -$, and the number of zeros of $\Delta\mathbb{E}CS_i(\theta)$ beyond the boundary is determined by the starting value: zero in the IR-constrained case and $\bar{\varepsilon}_i\,\mathrm{Cov}_i(R_\theta(\theta), R_{\mathcal{G}}(\theta))$ in the IR-unconstrained case, giving the patterns stated in the proposition. Under $\Delta\mathcal{G}_i(\theta) \leq 0$, all signs are reversed throughout. $\square$

C.7 Across-category allocation - Proof of Lemma 4

Proof. The feasible set

$$\mathcal{F} := \left\{ \{(k_i, I_i)\}_{i \in N} : \sum_{i \in N} \mu_i k_i = k,\ \sum_{i \in N} \mu_i I_i = I,\ k_i \geq 0,\ I_i \geq 0\ \forall i \right\}$$

is compact and convex. Continuity of each V_i follows from Berge's maximum theorem: the objective $\mathbb{E}_{(s,i)}[\Gamma_i(\theta) U(q,s)]$ and the constraints of the single-category problem are continuous in (k_i, I_i), and the feasible correspondence is continuous and compact-valued under the maintained assumptions. By Lemma 9 below, each V_i is concave. Hence the outer objective $W := \sum_i \mu_i V_i$ is continuous and concave on $\mathcal{F}$. Existence and convexity of the solution set follow from the extreme value theorem and concavity of W. Uniqueness under the stated condition follows from Lemma 9(i), which gives strict concavity of each V_i on the IR-constrained, capacity-binding region, hence strict concavity of W at the optimum.

Lemma 9 (Concavity of the single-category value function)**.** *For each category i, the value function $V_i(k_i, I_i)$ is concave on its feasible domain. More precisely:*

(i) on the IR-constrained, capacity-binding region, where $\beta_i(k_i, I_i) > 0$ and $\Pr(s \in T_i(k_i)) > 0$, the Hessian of V_i is negative definite and V_i is strictly concave;

(ii) when $\beta_i(k_i, I_i) = 0$, the Hessian is negative semidefinite;

(iii) when $T_i(k_i) = \varnothing$, the Hessian is negative semidefinite.

Proof of Lemma 9. Fix a category i. On binding states $s \in T_i(k_i)$, the pointwise first-order condition is

$$u(q_i(\theta, s), s)\, \mathcal{G}_i(\theta) = \varepsilon_i(s), \qquad \mathcal{G}_i(\theta) = \Gamma_i(\theta) + \beta_i J_i(\theta).$$

Step 1: k_i-derivative and the boundary term.

By the envelope theorem,

$$\frac{\partial V_i(k_i, I_i)}{\partial k_i} = \mathbb{E}_s \left[\varepsilon_i(s; k_i, I_i)\, \mathbf{1}_{\{s \in T_i(k_i)\}} \right].$$

Differentiating with respect to k_i gives a boundary term proportional to $\varepsilon_i(s_i)$. At the threshold state, the off-peak condition $u(q_i(\theta, s_i), s_i) = 0$ holds for every θ, and combined with the first-order condition this gives $\varepsilon_i(s_i) = 0$. Hence the boundary term vanishes and

$$\frac{\partial^2 V_i(k_i, I_i)}{\partial k_i^2} = \mathbb{E}_s \left[\frac{\partial \varepsilon_i(s; k_i, I_i)}{\partial k_i}\, \mathbf{1}_{\{s \in T_i(k_i)\}} \right].$$

Step 2: $\frac{\partial \varepsilon_i(s)}{\partial I_i}$ and the cross derivative.

Fix k_i and $s \in T_i$. Differentiating the first-order condition with respect to I_i and imposing the binding capacity constraint $\mathbb{E}_i[\partial_{I_i} q_i(\theta, s)] = 0$ yields $\frac{\partial \varepsilon_i(s)}{\partial I_i} = \partial_{I_i}\beta_i \, \mathbb{E}_i^w[uJ \mid s]$, Differentiating $\mathbb{E}\underline{CS}_i(k_i, \beta_i(k_i, I_i)) = 0$ with respect to I_i and using 4 gives $\partial_{I_i}\beta_i = 1/\mathbb{E}^w_{(s,i)}[\mathrm{Var}_i(uJ)]$ and hence $\frac{\partial \varepsilon_i(s)}{\partial I_i} = \mathbb{E}_i^w[uJ \mid s]/\mathbb{E}^w_{(s,i)}[\mathrm{Var}_i(uJ)]$. The boundary term vanishes again since $\varepsilon_i(s_i) = 0$, so

$$\frac{\partial V_i(k_i, I_i)}{\partial I_i} = -\tilde{\lambda}_i - \beta_i(k_i, I_i), \qquad \frac{\partial^2 V_i(k_i, I_i)}{\partial I_i^2} = -\frac{1}{\mathbb{E}^w_{(s,i)}[\mathrm{Var}_i(uJ)]} < 0, \qquad \frac{\partial^2 V_i(k_i, I_i)}{\partial k_i \partial I_i} = \frac{\mathbb{E}^w_{(s,i)}[uJ \mid s]}{\mathbb{E}^w_{(s,i)}[\mathrm{Var}_i(uJ)]}. \tag{10}$$

Step 3: $\frac{\partial^2 V_i(k_i, I_i)}{\partial k_i^2}$ on the IR-constrained region.

Holding β fixed gives $\frac{\partial \varepsilon_i(s)}{\partial k_i}|_\beta = -1/\overline{w}_i(s)$. Along the optimal path, the chain rule and 2 give

$$\frac{\partial \varepsilon_i(s)}{\partial k_i} = -\frac{1}{\overline{w}_i(s)} - \mathbb{E}_i^w[uJ \mid s] \frac{\mathbb{E}^w_{(s,i)}[uJ \mid s]}{\mathbb{E}^w_{(s,i)}[\mathrm{Var}_i(uJ)]}.$$

Integrating over binding states and defining $\Omega_i(k_i) := \mathbb{E}_s\left[\overline{w}_i(s)^{-1} \, \mathbf{1}_{\{s \in T_i(k_i)\}}\right] > 0$,

$$\frac{\partial^2 V_i(k_i, I_i)}{\partial k_i^2} = -\Omega_i(k_i) - \frac{\mathbb{E}^w_{(s,i)}[uJ \mid s]^2}{\mathbb{E}^w_{(s,i)}[\mathrm{Var}_i(uJ)]} < 0. \tag{11}$$

The Hessian on the IR-constrained, capacity-binding region is

$$\nabla^2 V_i(k_i, I_i) = \begin{pmatrix} -\Omega_i - \frac{\mathbb{E}^w_{(s,i)}[uJ|s]^2}{\mathbb{E}^w_{(s,i)}[\mathrm{Var}_i(uJ)]} & \frac{\mathbb{E}^w_{(s,i)}[uJ|s]}{\mathbb{E}^w_{(s,i)}[\mathrm{Var}_i(uJ)]} \\ \frac{\mathbb{E}^w_{(s,i)}[uJ|s]}{\mathbb{E}^w_{(s,i)}[\mathrm{Var}_i(uJ)]} & -\frac{1}{\mathbb{E}^w_{(s,i)}[\mathrm{Var}_i(uJ)]} \end{pmatrix},$$

with determinant $\det \nabla^2 V_i = \Omega_i \mathbb{E}^w_{(s,i)}[\mathrm{Var}_i(uJ)] > 0$. Both diagonal entries are negative and the determinant is positive, so $\nabla^2 V_i$ is negative definite, proving part (i).

Step 4: the remaining regimes.

If $\beta_i = 0$, then V_i is locally affine in I_i, so $\frac{\partial^2 V_i(k_i, I_i)}{\partial I_i^2} = 0$ and $\frac{\partial^2 V_i(k_i, I_i)}{\partial k_i \partial I_i} = 0$. If capacity binds on a positive-mass set of states, the fixed-β formula gives $\frac{\partial^2 V_i(k_i, I_i)}{\partial k_i^2} = -\Omega_i \leq 0$, so the Hessian is negative semidefinite, proving part (ii). If $T_i(k_i) = \varnothing$, all second derivatives vanish and the Hessian is the zero matrix, proving part (iii).

Step 5: global concavity across regime boundaries.

It remains to verify that ∇V_i is continuous across the two regime boundaries, which is required for the piecewise concavity argument to yield global concavity.

IR-constrained to IR-unconstrained. From Step 2, $\frac{\partial V_i(k_i,I_i)}{\partial I_i} = -\tilde{\lambda}_i - \beta_i$. As (k_i, I_i) approaches the boundary from the constrained side, $\beta_i \to 0^+$ by definition, so $\frac{\partial V_i(k_i,I_i)}{\partial I_i} \to -\tilde{\lambda}_i$, matching the unconstrained value. As $\beta_i \to 0^+$, the allocation and hence $\varepsilon_i(s)$ converge to their unconstrained counterparts, so $\frac{\partial V_i}{\partial k_i}$ is continuous across this boundary.

Capacity-binding to capacity-slack. As $s_i(k_i) \to \bar{s}$, $\mathbb{E}_s\left[\varepsilon_i(s)\,\mathbf{1}_{\{s\in T_i(k_i)\}}\right] \to 0$ continuously, matching the slack-side value of zero.

With ∇V_i continuous at both boundaries, fix any two feasible points x^0, x^1 and let $g_i(t) := V_i((1-t)x^0 + tx^1)$. Then g_i is piecewise C^2 with g_i' continuous at every switching point. On each regime piece $g_i''(t) \le 0$, so g_i' is nonincreasing on each piece. Continuity at switching points prevents upward jumps in g_i', so g_i' is globally nonincreasing and g_i is concave. Hence V_i is globally concave. □

□

C.8 Capacity ordering - Proof of Theorem 2

Lemma 10 (Revenue positivity)**.** *Under assumption Reg, if $I_i = 0$ then $\beta_i = 0$.*

Proof of Lemma 10. Suppose for contradiction that $I_i = 0$ and $\beta_i > 0$. Then complementary slackness on the IR/budget constraint requires $\mathbb{E}_{(s,i)}[U(q_i,s)J_i(\theta)] = 0$. However, since $J_i'(\theta) > 0$ and $q_i(\theta,s)$ is increasing in θ due to the IC constraint, hence both $J_i(\theta)$ and $U(q_i(\theta,s),s)$ are increasing in θ. Therefore $\mathrm{Cov}_i(J_i,U) \ge 0$, and combined with $\mathbb{E}_i[J_i] = \underline{\theta}_i > 0$ and $U > 0$ at any interior allocation, we obtain $\mathbb{E}_{(s,i)}[U(q_i,s)J_i(\theta)] > 0$ strictly, a contradiction. □

Proof. Since $V_i(k_i, I_i)$ is strictly concave in k_i by Lemma 9, the marginal scarcity rent

$$\Pi_i(k_i) := \frac{\partial V_i(k_i, I_i)}{\partial k_i} = \mathbb{E}_s\left[\varepsilon_i(s;k_i)\mathbf{1}_{\{s\in T_i(k_i)\}}\right]$$

is strictly decreasing in k_i. At the outer optimum, the stationarity condition equalizes $\Pi_i(k_i^*) = \zeta_k$ across all active categories. Therefore, to first order in $\Delta\lambda$, $k_j^* > k_i^* \iff \Delta\Pi_i(k_i^*) > 0$, where $\Delta\Pi_i(k_i^*) := \Pi_j(k_i^*) - \Pi_i(k_i^*)$ is the change in marginal scarcity rent induced by the preference perturbation, evaluated at the baseline capacity k_i^*. By the envelope theorem,

$$\Delta\Pi_i(k_i^*) = \mathbb{E}_s\left[\Delta\varepsilon_i(s;k_i^*)\,\mathbf{1}_{\{s\in T_i(k_i^*)\}}\right].$$

By Lemma 3, which holds for any allocation regime,

$$\frac{\Delta\varepsilon_i(s)}{\varepsilon_i(s)} = \mathbb{E}_i^w[R_{\mathcal{G}}(\theta) \mid s] \qquad \text{for each } s \in T_i(k_i^*),$$

where $R_{\mathcal{G}}(\theta) = \Delta\mathcal{G}_i(\theta)/\mathcal{G}_i(\theta)$ and $\mathcal{G}_i(\theta) = \Gamma_i(\theta) + J_i(\theta)\beta_i(k_i^*, I_i^*)$. Substituting gives

$$\Delta\Pi_i(k_i^*) = \mathbb{E}_s\left[\varepsilon_i(s;k_i^*)\,\mathbb{E}_i^w[R_{\mathcal{G}}(\theta) \mid s;k_i^*]\,\mathbf{1}_{\{s\in T_i(k_i^*)\}}\right],$$

which proves the first claim. Under Assumption CARA, $u_q = -\alpha u$, so $w_i(\theta,s) = 1/(\alpha\varepsilon_i(s))$ is independent of θ. Therefore $\mathbb{E}_i^w[R_{\mathcal{G}}(\theta) \mid s] = \mathbb{E}_i[R_{\mathcal{G}}(\theta)]$ for every state s, and

$$\Delta\Pi_i(k_i^*) = \Pi_i(k_i^*)\mathbb{E}_i[R_{\mathcal{G}}(\theta)]. \tag{12}$$

Since $\Pi_i(k_i^*) > 0$ at any interior optimum, the sign is determined solely by $\mathbb{E}_i[R_{\mathcal{G}}(\theta)]$. □

C.9 Preference shifts and effective social weights - Proof of Lemma 7

Proof. Write $\lambda_i^1(\theta) := \lambda_i(\theta)$, $\lambda_i^2(\theta) := \lambda_i^{\Delta}(\theta)$, and define

$$\tilde{\lambda}_i^k := \int_{\underline{\theta}_i}^{\overline{\theta}_i} \lambda_i^k(\theta)\,dG_i(\theta), \qquad \Lambda_i^k(\theta) := \frac{1}{g_i(\theta)}\int_{\theta}^{\overline{\theta}_i} \lambda_i^k(t)\,dG_i(t), \qquad k \in \{1,2\}.$$

Let $m(\theta) := \lambda_i^2(\theta)/\lambda_i^1(\theta)$, which is increasing in the upward-shift case and decreasing in the downward-shift case.

Core step: monotonicity of $\theta \mapsto \Lambda_i^2(\theta)/\Lambda_i^1(\theta)$. For each $\theta \in [\underline{\theta}_i, \overline{\theta}_i)$, define the probability measure Q_θ on $[\theta, \overline{\theta}_i]$ by

$$dQ_\theta(t) := \frac{\lambda_i^1(t)\,dG_i(t)}{\displaystyle\int_{\theta}^{\overline{\theta}_i} \lambda_i^1(x)\,dG_i(x)}\,\mathbf{1}\{t \geq \theta\}.$$

Since $\lambda_i^2(t) = m(t)\lambda_i^1(t)$, we have $\Lambda_i^2(\theta)/\Lambda_i^1(\theta) = \mathbb{E}_{Q_\theta}[m(\tilde{\theta})]$. For $\theta' < \theta$, Q_θ is $Q_{\theta'}$ conditioned on $\{\tilde{\theta} \geq \theta\}$. When m is increasing, this conditioning shifts mass toward larger values of m, so by first-order stochastic dominance $\mathbb{E}_{Q_\theta}[m(\tilde{\theta})] \geq \mathbb{E}_{Q_{\theta'}}[m(\tilde{\theta})]$. Hence $\theta \mapsto \Lambda_i^2(\theta)/\Lambda_i^1(\theta)$ is increasing on $[\underline{\theta}_i, \overline{\theta}_i)$. When m is decreasing the inequality reverses.

Proof of (i). Since $\theta \mapsto \Lambda_i^2(\theta)/\Lambda_i^1(\theta)$ is monotone, its logarithm is monotone too. Therefore, for almost every $\theta \in (\underline{\theta}_i, \overline{\theta}_i)$,

$$0 \leq \frac{d}{d\theta}\log\left(\frac{\Lambda_i^2(\theta)}{\Lambda_i^1(\theta)}\right) = \frac{\Lambda_i^{2\prime}(\theta)}{\Lambda_i^2(\theta)} - \frac{\Lambda_i^{1\prime}(\theta)}{\Lambda_i^1(\theta)},$$

in the upward case, which yields $\Lambda_i^{\Delta\prime}(\theta)/\Lambda_i^{\Delta}(\theta) \geq \Lambda_i'(\theta)/\Lambda_i(\theta)$ for almost every θ. The reverse inequality holds in the downward case.

Proof of (ii). Write

$$\mathcal{G}_i(\theta) = J_i(\theta)\big(\tilde{\lambda}_i + \beta_i\big) + \Lambda_i(\theta), \qquad \mathcal{G}_i^{\Delta}(\theta) = J_i(\theta)\big(\tilde{\lambda}_i^{\Delta} + \beta_i^{\Delta}\big) + \Lambda_i^{\Delta}(\theta).$$

Upward case. Assume m is increasing and $1 \leq c_{\mathcal{G}} \leq c_{\Lambda}$. The boundary evaluation gives

$$\frac{\Lambda_i^{\Delta}(\underline{\theta}_i)}{\Lambda_i(\underline{\theta}_i)} = \frac{\tilde{\lambda}_i^{\Delta}}{\tilde{\lambda}_i} = c_{\Lambda}.$$

Since the ratio is increasing by the core step and starts at $c_{\Lambda} \geq c_{\mathcal{G}}$,

$$\Lambda_i^{\Delta}(\theta) \geq c_{\Lambda}\, \Lambda_i(\theta) \geq c_{\mathcal{G}}\, \Lambda_i(\theta) \qquad \text{for all } \theta.$$

Using the definition of $c_{\mathcal{G}}$,

$$\mathcal{G}_i^{\Delta}(\theta) = c_{\mathcal{G}}\, J_i(\theta)\big(\tilde{\lambda}_i + \beta_i\big) + \Lambda_i^{\Delta}(\theta) \geq c_{\mathcal{G}}\, J_i(\theta)\big(\tilde{\lambda}_i + \beta_i\big) + c_{\mathcal{G}}\, \Lambda_i(\theta) = c_{\mathcal{G}}\, \mathcal{G}_i(\theta).$$

Since $c_{\mathcal{G}} \geq 1$ and $\mathcal{G}_i(\theta) > 0$, it follows that $\Delta\mathcal{G}_i(\theta) \geq (c_{\mathcal{G}} - 1)\mathcal{G}_i(\theta) \geq 0$ for all $\theta \in [\underline{\theta}_i, \overline{\theta}_i]$.

Downward case. Assume m is decreasing and $c_{\Lambda} \leq c_{\mathcal{G}} \leq 1$. The boundary evaluation gives $\Lambda_i^{\Delta}(\underline{\theta}_i)/\Lambda_i(\underline{\theta}_i) = c_{\Lambda} \leq c_{\mathcal{G}}$. Since the ratio is decreasing by the core step and starts at c_{Λ},

$$\Lambda_i^{\Delta}(\theta) \leq c_{\Lambda}\, \Lambda_i(\theta) \leq c_{\mathcal{G}}\, \Lambda_i(\theta) \qquad \text{for all } \theta.$$

Therefore $\mathcal{G}_i^{\Delta}(\theta) \leq c_{\mathcal{G}}\, \mathcal{G}_i(\theta)$, and since $c_{\mathcal{G}} \leq 1$ and $\mathcal{G}_i(\theta) > 0$, it follows that $\Delta\mathcal{G}_i(\theta) \leq (c_{\mathcal{G}} - 1)\mathcal{G}_i(\theta) \leq 0$ for all $\theta \in [\underline{\theta}_i, \overline{\theta}_i]$.

The special cases in the "in particular" clause follow immediately: if $\beta_i = \beta_i^{\Delta} = 0$ then $c_{\mathcal{G}} = c_{\Lambda} = \tilde{\lambda}_i^{\Delta}/\tilde{\lambda}_i$, so the ordering $c_{\mathcal{G}} \leq c_{\Lambda}$ holds with equality and the condition reduces to $\tilde{\lambda}_i^{\Delta} \geq \tilde{\lambda}_i$; if $\beta_i^{\Delta} + \tilde{\lambda}_i^{\Delta} = \beta_i + \tilde{\lambda}_i$ then $c_{\mathcal{G}} = 1$ and the condition reduces to $c_{\Lambda} = \tilde{\lambda}_i^{\Delta}/\tilde{\lambda}_i \geq 1$. □

C.10 Capacity ordering along an ordered preference path - Proof of Corollary 1

Proof. We proceed in three steps.

Step 1: Regime membership is monotone. Suppose for contradiction that there exists an adjacent pair $i-1, i$ such that $\beta_{i-1}^* = 0$ and $\beta_i^* > 0$. By Lemma 10, $\beta_i^* > 0$ implies $I_i^* > 0$. Hence the outer KKT condition for category i gives $\tilde{\lambda}_i + \beta_i^* = -\zeta_I$. For category $i-1$, the stationarity condition with $\beta_{i-1}^* = 0$ gives

$\tilde{\lambda}_{i-1} = -\zeta_I + \xi_{i-1}, \qquad \xi_{i-1} \geq 0$. Therefore $\tilde{\lambda}_{i-1} \geq -\zeta_I = \tilde{\lambda}_i + \beta_i^* > \tilde{\lambda}_i$, which contradicts $\tilde{\lambda}_i \geq \tilde{\lambda}_{i-1}$. Hence N^+ is an initial segment and N^0 a terminal segment of N.

Step 2: The sequence $(\beta_i^*)_{i \in N^+}$ is weakly decreasing. Fix $i \in N^+$. Since $\beta_i^* > 0$, Lemma 10 implies $I_i^* > 0$, so the outer KKT condition yields $\tilde{\lambda}_i + \beta_i^* = -\zeta_I$. Thus, for every $i \in N^+$, $\beta_i^* = -\zeta_I - \tilde{\lambda}_i$. Because $\tilde{\lambda}_i$ is weakly increasing in i under Assumption Order, it follows that β_i^* is weakly decreasing on N^+.

Step 3: Capacity is weakly increasing across adjacent categories. Fix an adjacent pair $i-1, i$, and let $t \longmapsto \lambda(t, \theta)$ be the ordered path from Assumption Order, with endpoints $\lambda(0, \theta) = \lambda_{i-1}(\theta), \qquad \lambda(1, \theta) = \lambda_i(\theta)$. Let $t \longmapsto \big(k^*(t), I^*(t), \beta^*(t)\big)$ denote the associated equilibrium path.

By the argument of Step 1 applied pointwise along the path, the regime $\beta^*(t)$ can move from > 0 to 0, but never from 0 to > 0. Hence, for almost every t, one of the following three local cases applies between t and $t + h$ for all sufficiently small $h > 0$:

$$(\beta^*(t), \beta^*(t+h)) = (0, 0), \qquad (\beta^*(t), \beta^*(t+h)) = (+, +), \qquad (\beta^*(t), \beta^*(t+h)) = (+, 0).$$

Under Assumption Order, the local perturbation from t to $t + h$ satisfies: $\frac{\lambda(t+h,\theta)}{\lambda(t,\theta)}$ is weakly increasing in θ, and $\tilde{\lambda}(t+h) \geq \tilde{\lambda}(t)$. We now verify that this implies $\mathcal{G}(\theta; t+h) - \mathcal{G}(\theta; t) \geq 0 \qquad \text{for all } \theta \in \Theta$.

Case 1: $\beta^*(t) = \beta^*(t+h) = 0$. Then $\mathcal{G}(\theta; t) = \Gamma(\theta; t), \qquad \mathcal{G}(\theta; t+h) = \Gamma(\theta; t+h)$, so $c_{\mathcal{G}} = c_\Lambda = \frac{\tilde{\lambda}(t+h)}{\tilde{\lambda}(t)} \geq 1$. Lemma 7 therefore yields $\mathcal{G}(\theta; t+h) - \mathcal{G}(\theta; t) \geq 0 \qquad \text{for all } \theta$.

Case 2: $\beta^*(t) > 0$ and $\beta^*(t+h) > 0$. Since both categories are in N^+, the outer KKT conditions give $\tilde{\lambda}(t) + \beta^*(t) = \tilde{\lambda}(t+h) + \beta^*(t+h) = -\zeta_I$. Hence $c_{\mathcal{G}} = 1, \qquad c_\Lambda = \frac{\tilde{\lambda}(t+h)}{\tilde{\lambda}(t)} \geq 1$. Lemma 7 again implies $\mathcal{G}(\theta; t+h) - \mathcal{G}(\theta; t) \geq 0 \qquad \text{for all } \theta$.

Case 3: $\beta^*(t) > 0$ and $\beta^*(t+h) = 0$. From the KKT conditions,

$$\tilde{\lambda}(t) + \beta^*(t) = -\zeta_I, \qquad \tilde{\lambda}(t+h) = -\zeta_I + \xi(t+h), \qquad \xi(t+h) \geq 0.$$

Therefore $\tilde{\lambda}(t+h) \geq \tilde{\lambda}(t) + \beta^*(t)$, so $c_{\mathcal{G}} = \frac{\tilde{\lambda}(t+h)}{\tilde{\lambda}(t)+\beta^*(t)} \geq 1$. Moreover,

$$c_\Lambda = \frac{\tilde{\lambda}(t+h)}{\tilde{\lambda}(t)} \geq \frac{\tilde{\lambda}(t+h)}{\tilde{\lambda}(t) + \beta^*(t)} = c_{\mathcal{G}}.$$

Lemma 7 then gives $\mathcal{G}(\theta; t+h) - \mathcal{G}(\theta; t) \geq 0 \qquad \text{for all } \theta$.

In all three cases, $\mathcal{G}(\theta; t+h) - \mathcal{G}(\theta; t) \geq 0 \qquad \text{for all } \theta$, hence $R_{\mathcal{G}}(\theta; t) := \frac{\partial_t \mathcal{G}(\theta;t)}{\mathcal{G}(\theta;t)} \geq 0 \qquad \text{for all } \theta$ for almost every t. Since the weights defining $\mathbb{E}^w[\cdot \mid s]$ are positive, $\mathbb{E}^w[R_{\mathcal{G}}(\theta) \mid s; t] \geq 0 \qquad \text{for every peak state } s$.

Applying Theorem 2 at the current optimum $\big(k^*(t), I^*(t)\big)$ yields $\frac{dk^*(t)}{dt} \geq 0$ for almost every $t \in [0,1]$. Integrating from $t = 0$ to $t = 1$ gives $k_i^* = k^*(1) \geq k^*(0) = k_{i-1}^*$.

Since this holds for every adjacent pair $i-1, i$, we conclude that $k_1^* \leq \cdots \leq k_n^*$. □

C.11 Revenue ordering - Proof of Proposition 2

Proof. Part (i). Assume both N^+ and N^0 nonempty. For $i \in N^+$, $\xi_i = 0$ and Lemma 4 gives $\tilde{\lambda}_i + \beta_i^* = -\zeta_I =: \tilde{\lambda}^*$. Since $\beta_i^* > 0$, $\tilde{\lambda}_i < \tilde{\lambda}^*$. For $j \in N^0$, $\beta_j^* = 0$, so $\tilde{\lambda}_j = \tilde{\lambda}^* + \xi_j \geq \tilde{\lambda}^*$. Hence $\max_{N^+} \tilde{\lambda}_i < \tilde{\lambda}^* \leq \min_{N^0} \tilde{\lambda}_j$.

Part (ii). General case. Since $\beta_i^*, \beta_j^* > 0$, Lemma 4 gives $\tilde{\lambda}_i + \beta_i^* = \tilde{\lambda}_j + \beta_j^*$, so $\Delta\beta^* = -\Delta\tilde{\lambda}$. Let $\Delta k := k_j^* - k_i^* > 0$ and $\Delta I := I_j^* - I_i^*$. A first-order expansion of $\beta_j^* = \beta_j(k_j^*, I_j^*; \lambda_j)$ around $(k_i^*, I_i^*; \lambda_i)$ gives

$$\frac{\partial \beta_i}{\partial k_i}\Delta k + \frac{\partial \beta_i}{\partial I_i}\Delta I + \Delta_\lambda \beta_i = -\Delta\tilde{\lambda},$$

where $\Delta_\lambda \beta_i$ is the direct effect of the preference perturbation at fixed (k_i^*, I_i^*). It is obtained from the implicit-function derivative of the binding IR constraint $F_i(\beta_i; k, I, \lambda) := \mathbb{E}\underline{CS}_i(k, I; \beta_i, \lambda) = 0$, holding β_i fixed in the numerator:

$$\Delta_\lambda \beta_i = -\frac{\partial_\lambda F_i|_{\beta_i \text{ fixed}}}{\partial_{\beta_i} F_i} = -\frac{\mathcal{C}_i^0}{\mathbb{E}^w_{(s,i)}[\mathrm{Var}_i(uJ)]},$$

with

$$\mathcal{C}_i^0 := \mathbb{E}_{(s,i)}\Big[\varepsilon_i(s) R_\theta(\theta) R_q(\theta) \Big(R_{\mathcal{G}}^0(\theta) - \mathbb{E}_i^w[R_{\mathcal{G}}^0(\theta) \mid s]\Big) \mathbf{1}_{\{s \in T_i\}}\Big], \qquad R_{\mathcal{G}}^0(\theta) := \frac{J_i(\theta)\Delta\tilde{\lambda} + \Delta\Lambda_i(\theta)}{\mathcal{G}_i(\theta)} = R_\theta(\theta)\Delta\tilde{\lambda} + R_{\mathcal{G}}(\theta).$$

Using $\partial_{I_i}\beta_i = 1/\mathbb{E}^w_{(s,i)}[\mathrm{Var}_i(uJ)] > 0$ from Lemma 9, we obtain

$$\Delta I = -\frac{\Delta\tilde{\lambda} + \frac{\partial \beta_i}{\partial k_i}\Delta k + \Delta_\lambda \beta_i}{\frac{\partial \beta_i}{\partial I_i}},$$

so $\Delta I < 0$ iff $\Delta\tilde{\lambda} + \frac{\partial \beta_i}{\partial k_i}\Delta k + \Delta_\lambda \beta_i > 0$. The sign of $\partial\beta_i/\partial k_i$ is opposite to the sign of EuJ by 5.

CARA reduction. Under Assumption CARA, $w_i = 1/(\alpha\varepsilon_i(s))$ is θ-independent, and

$$\mathbb{E}^w_{(s,i)}[uJ \mid s] = \bar{\varepsilon}_i \mathbb{E}_i[R_\theta(\theta)], \tag{13}$$

$$\mathbb{E}^w_{(s,i)}[\mathrm{Var}_i(uJ)] = \frac{\bar{\varepsilon}_i}{\alpha}\mathrm{Var}_i(R_\theta(\theta)), \tag{14}$$

so $\partial_{k_i}\beta_i = -\alpha\mathbb{E}_i[R_\theta(\theta)]/\mathrm{Var}_i(R_\theta(\theta))$. By 12, $\Delta\Pi_i = \Pi_i\mathbb{E}_i[R_\mathcal{G}(\theta)]$, and 11 gives $\Pi_i'(k_i^*) = -\alpha\Pi_i\mathbb{E}_i[R_\theta(\theta)^2]/\mathrm{Var}_i(R_\theta(\theta))$. Hence $\Delta k = -\Delta\Pi_i/\Pi_i' = \mathbb{E}_i[R_\mathcal{G}(\theta)]\mathrm{Var}_i(R_\theta(\theta))/(\alpha\mathbb{E}_i[R_\theta(\theta)^2])$, and

$$\Delta k\,\frac{\partial\beta_i}{\partial k_i} = -\frac{\mathbb{E}_i[R_\mathcal{G}(\theta)]\mathbb{E}_i[R_\theta(\theta)]}{\mathbb{E}_i[R_\theta(\theta)^2]}.$$

Substituting $R^0_\mathcal{G}(\theta) = R_\theta(\theta)\Delta\tilde{\lambda} + R_\mathcal{G}(\theta)$ into $\mathcal{C}_i^0$ and using $\mathrm{Cov}_i(R_\theta(\theta), R_\theta(\theta)) = \mathrm{Var}_i(R_\theta(\theta))$:

$$\mathcal{C}_i^0 = \Delta_\lambda\mathbb{E}_{(s,i)}[U(q_i(\theta,s),s)J_i(\theta)] = \frac{\bar{\varepsilon}_i}{\alpha}\mathrm{Cov}_i(R_\theta(\theta), R^0_\mathcal{G}(\theta)). \tag{15}$$

Using (14), the $\bar{\varepsilon}_i/\alpha$ factors cancel and

$$\Delta_\lambda\beta_i = -\frac{\mathrm{Cov}_i(R_\theta(\theta), R^0_\mathcal{G}(\theta))}{\mathrm{Var}_i(R_\theta(\theta))} = -\Delta\tilde{\lambda} - \frac{\mathrm{Cov}_i(R_\theta(\theta), R_\mathcal{G}(\theta))}{\mathrm{Var}_i(R_\theta(\theta))}. \tag{16}$$

Plugging into the condition $\Delta\tilde{\lambda} + \Delta k\,\partial_k\beta_i + \Delta_\lambda\beta_i > 0$ and cancelling $\Delta\tilde{\lambda}$ yields $\Delta I < 0$ iff

$$\frac{\mathbb{E}_i[R_\mathcal{G}(\theta)]\mathbb{E}_i[R_\theta(\theta)]}{\mathbb{E}_i[R_\theta(\theta)^2]} + \frac{\mathrm{Cov}_i(R_\theta(\theta), R_\mathcal{G}(\theta))}{\mathrm{Var}_i(R_\theta(\theta))} < 0. \qquad \square$$

C.12 Optimal investment - Proof of Lemma 5

Proof. Define the outer indirect value $\overline{W}(K, I) := \max_{\{k_i, I_i\}_{i\in N}} \sum_{i\in N} \mu_i V_i(k_i, I_i)$ subject to

$$\sum_{i\in N}\mu_i k_i = K, \qquad \sum_{i\in N}\mu_i I_i = I, \qquad k_i \geq 0, \qquad I_i \geq 0 \quad \forall i \in N.$$

By Lemma 4, $\overline{W}$ is well defined and continuous, and by Lemma 9 it is concave in (K, I). Moreover, by 10, $\frac{\partial V_i(k_i,I_i)}{\partial I_i} = -\tilde{\lambda}_i - \beta_i(k_i, I_i) < 0$, the function $\overline{W}(K, I)$ is strictly decreasing in I.

Now define $W(k) := \overline{W}(k, I(k))$. Let $k_0, k_1 \geq 0$ and $t \in (0, 1)$. By concavity of $\overline{W}$ and convexity of I,

$$\begin{aligned} W(tk_0 + (1-t)k_1) &= \overline{W}(tk_0 + (1-t)k_1,\ I(tk_0 + (1-t)k_1)) \\ &\geq \overline{W}(tk_0 + (1-t)k_1,\ tI(k_0) + (1-t)I(k_1)) \\ &\geq t\,W(k_0) + (1-t)\,W(k_1). \end{aligned}$$

Since I is strictly convex and $\overline{W}$ is strictly decreasing in its second argument, the first inequality is strict whenever $k_0 \neq k_1$. Hence W is strictly concave, so any maximizer is unique.

By the maintained Inada condition on u, the marginal value of capacity tends to $+\infty$ as $k \downarrow 0$, while strict convexity of I implies $W(k) \to -\infty$ for large k. Therefore the unique maximizer satisfies $k^* > 0$.

At any point of differentiability, the envelope theorem for the outer problem gives $W'(k) = \zeta_k + \zeta_I I'(k)$, where ζ_k and ζ_I are the multipliers on the two aggregate constraints. Since k^* is an interior maximizer, $0 = W'(k^*) = \zeta_k + \zeta_I I'(k^*)$. Let $\tilde{\lambda}^* := -\zeta_I > 0$. Then $\tilde{\lambda}^* I'(k^*) = \zeta_k$.

By Lemma 4, $\Pi_i(k_i^*, I_i^*) = \zeta_k$ for every i, where, by the single-category envelope formula and the first-order condition FOC_q,

$$\Pi_i(k_i^*, I_i^*) = \mathbb{E}_{(s,i)}\left[u(q_i(\theta, s), s)\, \mathcal{G}_i(\theta)\, \mathbf{1}_{\{s \in T_i\}}\right].$$

Let $\mu^+ := \sum_{i \in N^+} \mu_i$. Since $\Pi_i(k_i^*, I_i^*) = \zeta_k$ for every contributing category, $\mu^+ \zeta_k = \sum_{i \in N^+} \mu_i \Pi_i(k_i^*, I_i^*)$. Combining this with $\tilde{\lambda}^* I'(k^*) = \zeta_k$ yields

$$I'(k^*) = \frac{1}{\mu^+} \sum_{i \in N^+} \mu_i\, \mathbb{E}_{(s,i)}\left[u(q_i(\theta, s), s)\, \frac{\mathcal{G}_i(\theta)}{\tilde{\lambda}^*}\, \mathbf{1}_{\{s \in T_i\}}\right].$$

For the comparative statics below, it is convenient to note the equivalent whole-N representation $I'(k^*) = \frac{1}{\tilde{\lambda}^*} \sum_{i \in N} \mu_i \Pi_i(k_i^*, I_i^*)$, which follows from the same identity $\Pi_i(k_i^*, I_i^*) = \zeta_k$.

Finally, the outer KKT condition is

$$\tilde{\lambda}_i + \beta_i(k_i^*, I_i^*) = \tilde{\lambda}^* + \xi_i, \qquad \xi_i I_i^* = 0, \qquad \xi_i \geq 0.$$

If $i \in N^+$, then $\beta_i^* > 0$, so $\tilde{\lambda}_i < \tilde{\lambda}_i + \beta_i^* = \tilde{\lambda}^* + \xi_i$. Since contributing categories are interior in the revenue margin, $\xi_i = 0$, and therefore $\tilde{\lambda}_i < \tilde{\lambda}^*$. If instead $i \in N^0$, then $\beta_i^* = 0$, so $\tilde{\lambda}_i = \tilde{\lambda}^* + \xi_i \geq \tilde{\lambda}^*$. Thus $\max_{i \in N^+} \tilde{\lambda}_i < \tilde{\lambda}^* \leq \min_{i \in N^0} \tilde{\lambda}_i$. $\square$

C.13 Investment and preference perturbations - Proof of Proposition 3

Proof. Let Δ_λ denote the first-order effect of the perturbation of $\lambda_i(\theta)$ evaluated at the baseline equilibrium allocation.

Differentiating the optimality condition $W'(k^*) = 0$ with respect to the perturbation parameter. By the envelope theorem applied to the outer problem, $W'(k) = \zeta_k + \zeta_I I'(k)$, so $W''(k^*)\, \Delta k^* = -\Delta_\lambda \zeta_k - I'(k^*)\, \Delta_\lambda \zeta_I$. From Lemma 4, $\zeta_k = \sum_{i \in N} \mu_i \Pi_i$ at the optimum, so $\Delta_\lambda \zeta_k = \sum_{i \in N} \mu_i \Delta_\lambda \Pi_i$. Since $\zeta_I = -\tilde{\lambda}^*$, we have $\Delta_\lambda \zeta_I = -\Delta \tilde{\lambda}^*$. Substituting,

$$W''(k^*)\, \Delta k^* = I'(k^*)\, \Delta \tilde{\lambda}^* - \sum_{i \in N} \mu_i\, \Delta_\lambda \Pi_i.$$

Under Assumption CARA, 12 gives, for every category i, $\Delta_\lambda \Pi_i = \Pi_i \, \mathbb{E}_i[R_{\mathcal{G},i}]$. Substituting into the previous display gives

$$W''(k^*)\,\Delta k^* = I'(k^*)\,\Delta\tilde{\lambda}^* - \sum_{i\in N} \mu_i \, \Pi_i \, \mathbb{E}_i[R_{\mathcal{G},i}].$$

It remains to characterize $\Delta\tilde{\lambda}^*$. For every contributing category $i \in N^+$, the revenue constraint binds. Since the perturbation is mean-preserving, $\Delta\tilde{\lambda}_i = 0$, and Lemma 5 implies $\tilde{\lambda}_i + \beta_i^* = \tilde{\lambda}^* \quad$ for all $i \in N^+$, hence $\Delta\beta_i^* = \Delta\tilde{\lambda}^*$. Taking the first-order effect of the perturbation in the binding condition gives

$$\mathbb{E}^{w}_{(s,i)}[\mathrm{Var}_i(uJ)](k_i^*)\,\Delta\tilde{\lambda}^* + \Delta_\lambda \mathbb{E}_{(s,i)}[U(q_i(\theta,s),s)\,J_i(\theta)] = 0.$$

Under Assumption CARA, mean preserving assumption, 14 and 15 imply

$$\mathbb{E}^{w}_{(s,i)}[\mathrm{Var}_i(uJ)](k_i^*) = \frac{\bar{\varepsilon}_i}{\alpha}\,\mathrm{Var}_i(R_{\theta,i}(\theta))\,, \qquad \Delta_\lambda \mathbb{E}_{(s,i)}[U(q_i(\theta,s),s)\,J_i(\theta)] = \frac{\bar{\varepsilon}_i}{\alpha}\,\mathrm{Cov}_i(R_{\theta,i}(\theta), R_{\mathcal{G},i}(\theta))\,,$$

where $\bar{\varepsilon}_i := \mathbb{E}_s\left[\varepsilon_i(s)\mathbf{1}_{\{s\in T_i\}}\right]$. Therefore, summing over all contributing categories $i \in N^+$ with weights μ_i, we obtain

$$\Delta\tilde{\lambda}^* = -\frac{\sum_{i\in N^+} \mu_i \, \bar{\varepsilon}_i \, \mathrm{Cov}_i(R_{\theta,i}(\theta), R_{\mathcal{G},i}(\theta))}{\sum_{i\in N^+} \mu_i \, \bar{\varepsilon}_i \, \mathrm{Var}_i(R_{\theta,i}(\theta))}.$$

Finally, under a smooth ordered path that tilts redistributive preferences toward higher types, we have $\mathbb{E}_i[R_{\mathcal{G},i}] > 0$ for all $i \in N$. Hence the average effect $-\sum_{i\in N} \mu_i \, \Pi_i \, \mathbb{E}_i[R_{\mathcal{G},i}]$ is negative. Since $W''(k^*) < 0$, this channel pushes Δk^* upward. The overall sign of Δk^* nevertheless also depends on the induced change in $\tilde{\lambda}^*$. $\square$

C.14 Implementation feasibility - Proof of Proposition 4

Proof. By 1, $\frac{\partial \mathbb{E}\underline{CS}_i}{\partial k_i} = \mathbb{E}^{w}_{(s,i)}[uJ \mid s]$. Under the conditions of Lemma 6, the map $k_i \mapsto \mathbb{E}^{w}_{(s,i)}[uJ \mid s]$ is single-crossing. Hence $\frac{\partial \mathbb{E}\underline{CS}_i}{\partial k_i}$ is monotone and crosses zero at most once in a given direction.

When EuJ crosses from $+$ to $-$, $\frac{\partial \mathbb{E}\underline{CS}_i}{\partial k_i}$ is decreasing and crosses zero at most once from above, so $\mathbb{E}\underline{CS}_i$ is single-peaked on $[0, k_i^+]$. When EuJ crosses from $-$ to $+$, $\frac{\partial \mathbb{E}\underline{CS}_i}{\partial k_i}$ is increasing and $\mathbb{E}\underline{CS}_i$ is single-dipped. At $k_i = 0$ the capacity constraint binds at zero, so $q_i(\theta,s) = 0$ for all (θ,s), giving $\mathbb{E}\underline{CS}_i(0) = -I_i < 0 \quad$ for any $I_i > 0$.

Case: EuJ crosses from $+$ to $-$. Suppose first that $\mathbb{E}\underline{CS}_i(k_i^+) \geq 0$. Since $\mathbb{E}\underline{CS}_i(0) < 0$ and $\mathbb{E}\underline{CS}_i(k_i^+) \geq 0$, continuity gives at least one zero in $(0, k_i^+]$. Single-peakedness implies that this zero is unique, defining $\tilde{k}_i$, and $\mathbb{E}\underline{CS}_i(k_i) \geq 0 \iff k_i \in [\tilde{k}_i, k_i^+]$.

This proves part (i).

Suppose instead that $\mathbb{E}\underline{CS}_i(k_i^+) < 0$. Then $\mathbb{E}\underline{CS}_i$ is negative at both endpoints. If the peak value is negative, then $\mathbb{E}\underline{CS}_i < 0$ throughout. If the peak value is nonnegative, continuity and single-peakedness yield exactly two zeros $\tilde{k}_i^- < \tilde{k}_i^+$ in $(0, k_i^+)$ such that $\mathbb{E}\underline{CS}_i(k_i) \geq 0 \iff k_i \in [\tilde{k}_i^-, \tilde{k}_i^+]$. This proves part (ii).

Case: EuJ crosses from $-$ to $+$. Then $\mathbb{E}\underline{CS}_i$ is single-dipped with $\mathbb{E}\underline{CS}_i(0) < 0$. Either the trough value is negative and $\mathbb{E}\underline{CS}_i < 0$ throughout, or $\mathbb{E}\underline{CS}_i$ rises above zero after the trough, crossing zero exactly once from below at a unique $\tilde{k}_i \in (0, k_i^+]$, and $\mathbb{E}\underline{CS}_i(k_i) \geq 0 \iff k_i \in [\tilde{k}_i, k_i^+]$. $\square$

C.15 Non-monotone reallocation under preference perturbations - Proof of Proposition 5

Proof. Fix a peak state $s \in T_i$. The proof is based on the reallocation identity from Lemma 3:

$$\Delta q_i(\theta, s) = R_q(\theta)\left(R_{\mathcal{G}}(\theta) - \mathbb{E}_i^w[R_{\mathcal{G}}(\theta) \mid s]\right), \tag{17}$$

together with the fact that, at fixed capacity,

$$\mathbb{E}_i[\Delta q_i(\theta, s)] = 0. \tag{18}$$

Indeed, both the baseline and perturbed allocations satisfy the same state-by-state capacity constraint in state s.

We proceed in three steps.

Step 1: endpoint values of $R_{\mathcal{G}}(\theta)$.

Under the IR-unconstrained allocation, $\beta_i = 0$, using $J_i(\underline{\theta}_i) = \underline{\theta}_i - \gamma_i(\underline{\theta}_i)$ and $\Lambda_i(\underline{\theta}_i) = \gamma_i(\underline{\theta}_i)\tilde{\lambda}_i$, we obtain $\mathcal{G}_i(\underline{\theta}_i) = \underline{\theta}_i\tilde{\lambda}_i$. Likewise, since $\Lambda_i(\overline{\theta}_i) = 0$ and $J_i(\overline{\theta}_i) = \overline{\theta}_i$, we obtain $\mathcal{G}_i(\overline{\theta}_i) = \overline{\theta}_i\tilde{\lambda}_i$. Therefore,

$$R_{\mathcal{G}}(\underline{\theta}_i) = R_{\mathcal{G}}(\overline{\theta}_i) = \frac{\Delta\tilde{\lambda}_i}{\tilde{\lambda}_i}. \tag{19}$$

Under the IR-constrained allocation, assume $\beta_i^\Delta + \tilde{\lambda}_i^\Delta = \beta_i + \tilde{\lambda}_i$, that is, $\Delta\beta_i + \Delta\tilde{\lambda}_i = 0$. it follows that

$$\Delta\mathcal{G}_i(\theta) = (\Delta\tilde{\lambda}_i + \Delta\beta_i)J_i(\theta) + \Delta\Lambda_i(\theta) = \Delta\Lambda_i(\theta).$$

Evaluating at the upper endpoint gives $\Delta\mathcal{G}_i(\overline{\theta}_i) = \Delta\Lambda_i(\overline{\theta}_i) = 0$, since $\Lambda_i(\overline{\theta}_i) = 0$. Hence

$$R_{\mathcal{G}}(\overline{\theta}_i) = 0. \tag{20}$$

At the lower endpoint,

$$\Delta\mathcal{G}_i(\underline{\theta}_i) = \Delta\Lambda_i(\underline{\theta}_i) = \gamma_i(\underline{\theta}_i)\Delta\tilde{\lambda}_i,$$

so

$$R_{\mathcal{G}}(\underline{\theta}_i) = \frac{\gamma_i(\underline{\theta}_i)}{\mathcal{G}_i(\underline{\theta}_i)}\Delta\tilde{\lambda}_i. \tag{21}$$

In particular, the lower-endpoint value is of order $\Delta\tilde{\lambda}_i$, whereas the upper-endpoint value is exactly zero. Under a mean-preserving perturbation, $\Delta\tilde{\lambda}_i = 0$, so both endpoint values are zero.

Step 2: proof of part (i).

We first consider the IR-unconstrained allocation. By 19, the bracket in 17 takes the same value at both endpoints. Hence

$$\Delta q_i(\underline{\theta}_i, s) = R_q(\underline{\theta}_i)\left(\frac{\Delta\tilde{\lambda}_i}{\tilde{\lambda}_i} - \mathbb{E}_i^w[R_{\mathcal{G}}(\theta) \mid s]\right),$$

$$\Delta q_i(\overline{\theta}_i, s) = R_q(\overline{\theta}_i)\left(\frac{\Delta\tilde{\lambda}_i}{\tilde{\lambda}_i} - \mathbb{E}_i^w[R_{\mathcal{G}}(\theta) \mid s]\right).$$

Since $R_q(\theta) > 0$, the two endpoint values of $\Delta q_i(\cdot, s)$ have the same sign.

Moreover, because $R_{\mathcal{G}}(\theta)$ is not a.e. constant, the term $R_{\mathcal{G}}(\theta) - \mathbb{E}_i^w[R_{\mathcal{G}}(\theta) \mid s]$ is not a.e. zero, so 17 implies that $\Delta q_i(\cdot, s)$ is not identically zero. Combining this with 18, $\Delta q_i(\cdot, s)$ must take both positive and negative values. A function on an interval that takes both signs and whose two endpoint values have the same sign cannot be monotone. This proves the first claim in part (i).

Consider next the IR-constrained allocation. By 20–21, $\Delta q_i(\overline{\theta}_i, s) = -R_q(\overline{\theta}_i)\mathbb{E}_i^w[R_{\mathcal{G}}(\theta) \mid s]$, whereas

$$\Delta q_i(\underline{\theta}_i, s) = R_q(\underline{\theta}_i)\left(\frac{\gamma_i(\underline{\theta}_i)}{\mathcal{G}_i(\underline{\theta}_i)}\Delta\tilde{\lambda}_i - \mathbb{E}_i^w[R_{\mathcal{G}}(\theta) \mid s]\right).$$

Hence the two endpoint values have the same sign whenever

$$\left|\frac{\gamma_i(\underline{\theta}_i)}{\mathcal{G}_i(\underline{\theta}_i)}\Delta\tilde{\lambda}_i\right| < \left|\mathbb{E}_i^w[R_{\mathcal{G}}(\theta) \mid s]\right|.$$

Therefore, for $|\Delta\tilde{\lambda}_i|$ sufficiently small, the two endpoints inherit the same sign. Since $\Delta q_i(\cdot, s)$ is again not identically zero and satisfies 18, it cannot be monotone. This proves the second claim in part (i).

Step 3: proof of part (ii).

Assume now that the perturbation is mean-preserving, so $\Delta\tilde{\lambda}_i = 0$. Under the maintained restriction $\Delta\beta_i + \Delta\tilde{\lambda}_i = 0$, this also gives $\Delta\beta_i = 0$. By Step 1, $R_{\mathcal{G}}(\underline{\theta}_i) = R_{\mathcal{G}}(\overline{\theta}_i) = 0$. Hence, evaluating 17 at the two endpoints,

$$\Delta q_i(\underline{\theta}_i, s) = -R_q(\underline{\theta}_i)\mathbb{E}_i^w[R_{\mathcal{G}}(\theta) \mid s], \qquad \Delta q_i(\overline{\theta}_i, s) = -R_q(\overline{\theta}_i)\mathbb{E}_i^w[R_{\mathcal{G}}(\theta) \mid s]. \tag{22}$$

Thus the two endpoints have the same sign, and by part (i), $\Delta q_i(\cdot, s)$ is not monotone.

Because $\Delta q_i(\cdot, s)$ is continuous, not identically zero, and satisfies 18, it must cross zero at least twice. Therefore the cutoff set $\{\theta : \Delta q_i(\theta, s) = 0\}$ contains at least two elements. Let $\theta_- \leq \theta_+$ denote its minimum and maximum elements.

Suppose first that $\Delta \mathcal{G}_i(\theta) \geq 0$ for all θ. Since $\mathcal{G}_i(\theta) > 0$, this implies

$$R_{\mathcal{G}}(\theta) = \frac{\Delta \mathcal{G}_i(\theta)}{\mathcal{G}_i(\theta)} \geq 0 \qquad \text{for all } \theta,$$

and therefore $\mathbb{E}_i^w[R_{\mathcal{G}}(\theta) \mid s] \geq 0$. By 22, both endpoint values of $\Delta q_i(\cdot, s)$ are nonpositive. Since $\Delta q_i(\cdot, s)$ has zero mean and is not identically zero, it must be strictly positive on an interior subset of positive measure. It follows that

$$\Delta q_i(\theta, s) \leq 0 \qquad \text{for } \theta \in [\underline{\theta}_i, \theta_-) \cup (\theta_+, \overline{\theta}_i].$$

The case $\Delta \mathcal{G}_i(\theta) \leq 0$ is symmetric. In that case, $R_{\mathcal{G}}(\theta) \leq 0 \qquad$ for all θ, so $\mathbb{E}_i^w[R_{\mathcal{G}}(\theta) \mid s] \leq 0$, and 22 implies that both endpoint values of $\Delta q_i(\cdot, s)$ are nonnegative. Hence

$$\Delta q_i(\theta, s) \geq 0 \qquad \text{for } \theta \in [\underline{\theta}_i, \theta_-) \cup (\theta_+, \overline{\theta}_i].$$

This proves part (ii). $\square$

C.16 Marginal response of category allocations to aggregate capacity - Proof of Lemma 8

Proof. Contributing categories ($i \in N^+$). Both outer first-order conditions hold with equality and dI_i^* is free. Differentiating $\frac{\partial V_i}{\partial k_i} = \tilde{\lambda}^* I'(k^*)$ and $\frac{\partial V_i}{\partial I_i} = -\tilde{\lambda}^*$ with respect to k and using 10–11 gives

$$dI_i^* = \mathbb{E}_{(s,i)}^w[uJ \mid s]\, dk_i^* + \mathbb{E}_{(s,i)}^w[\mathrm{Var}_i(uJ)]\, d\tilde{\lambda}^*.$$

Substituting into the differentiated capacity condition and using $-\frac{\partial^2 V_i}{\partial k_i^2} - (\mathbb{E}_{(s,i)}^w[uJ \mid s])^2 / \mathbb{E}_{(s,i)}^w[\mathrm{Var}_i(uJ)] = \Omega_i$ yields

$$\Omega_i\, dk_i^* = \Big(\mathbb{E}_{(s,i)}^w[uJ \mid s] - I'(k^*)\Big)\, d\tilde{\lambda}^* - \tilde{\lambda}^* I''(k^*)\, dk.$$

Non-contributing categories ($i \in N^0$). Since $I_i^* = 0$ identically on the neighborhood, $dI_i^* = 0$. By Lemma 9(ii), $\beta_i = 0$ implies $\frac{\partial^2 V_i}{\partial k_i \partial I_i} = 0$ and $\frac{\partial^2 V_i}{\partial k_i^2} = -\Omega_i$. Differentiating only the capacity condition gives $\Omega_i\, dk_i^* = -I'(k^*)\, d\tilde{\lambda}^* - \tilde{\lambda}^* I''(k^*)\, dk$. This coincides with the contributing formula with $\mathbb{E}_{(s,i)}^w[uJ \mid s]$ replaced by $\hat{A}_i = 0$, so both regimes are covered by

$$\Omega_i \, dk_i^* = \big(\hat{A}_i - I'(k^*)\big) \, d\tilde{\lambda}^* - \tilde{\lambda}^* I''(k^*) \, dk.$$

Solving for $d\tilde{\lambda}^*$. Substituting into the aggregate feasibility condition $\sum_{j \in N} \mu_j \, dk_j^* = dk$ and rearranging,

$$d\tilde{\lambda}^* \left(\sum_{j \in N^+} \mu_j \Omega_j^{-1} \mathbb{E}^w_{(s,j)}[uJ \mid s] - I'(k^*) \sum_{j \in N} \mu_j \Omega_j^{-1} \right) = \mathcal{B} \, dk,$$

where $\mathcal{B} = 1 + \tilde{\lambda}^* I''(k^*) \sum_{j \in N} \mu_j \Omega_j^{-1} > 0$. Substituting back yields the formula in the lemma. □