

Sorting on Survival Gains Under Cost Exposure

Discussant comments for AHEW

Santiago Acerenza Fleitas

Universidad ORT Uruguay

Central theme: cost exposure as targeting device or barrier to high-value care.

- ➊ **Set up the economic problem:** cost exposure can either target scarce private capacity or screen out high-gain patients.
- ➋ **Reconstruct the paper's logic:** Chilean institutions → Roy choice → resistance → MTR/MTE curves.
- ➌ **Develop three discussion points:**
 - the economic meaning of the interaction instrument;
 - monotonicity with both main effects controlled;
 - repeated hospitalizations, the estimand, and inference.
- ➍ **Return to the evidence:** interpret the curves, treatment parameters, mechanism, and policy implications in light of those points.

The paper in one slide

- **Question:** do patients sort into private care based on expected survival gains when care involves cost exposure?
- **Setting:** publicly insured surgical patients in Chile's mixed public-private system.
- **Treatment:** private rather than public hospital admission.
- **Method:** MTE framework for binary mortality using continuous surgery-year variation generated by reimbursement, provider availability, and their interaction.
- **Main finding:** mortality gains become larger with resistance to private care; the pattern is more pronounced among lower-income patients.

My discussion in three points

- 1 **Economic interpretation.** When does cost exposure target high-gain patients? Through which concrete channel does $R_{st}N_{st}$ change effective financial exposure or effective access in the Roy index?
- 2 **Monotonicity.** What economic structure makes the interaction shift private-care choice in a common direction for every latent type when both main effects are controlled?
- 3 **Unit of latent sorting.** If the same patient can contribute several diagnosis-specific admissions, does the estimated MTE rank persistent patient type θ_i , episode-specific resistance, or a mixture of both?

Roy interpretation $\longrightarrow$ monotonicity $\longrightarrow$ patient type versus hospitalization margin.

The opening conundrum

- Public care is free at the point of use but constrained by capacity and waiting times.
- Private care may provide faster access, but publicly insured patients face financial exposure.
- If patients with the largest survival gains can bear that cost, price exposure may target scarce private capacity.
- If high-gain patients are also the most financially constrained, the same mechanism can screen them out.

Does cost exposure improve allocation, or create misallocation?

Chile: why the application is relevant

- Public care is generally free at point of use, but capacity constrained and linked to referrals and residence.
- Private care may provide faster access and greater choice, but publicly insured patients face financial exposure.
- FONASA and the MLE voucher scheme create a cost-sharing margin, while private hospitals decide whether and which vouchers to accept.

Treatment, outcomes, and sign convention

- $D_i = 1$ denotes private hospital admission and $D_i = 0$ public hospital admission.
- For $h \in \{7, 28, 90, 180\}$, define $M_i^h(d) = 1$ if the patient would die by horizon h under sector d .
- Observed mortality is

$$M_i^h = D_i M_i^h(1) + (1 - D_i) M_i^h(0).$$

- The mortality effect and corresponding survival gain are

$$\tau_i^h = M_i^h(1) - M_i^h(0), \quad G_i^h = M_i^h(0) - M_i^h(1) = -\tau_i^h.$$

$$\tau_i^h < 0 \iff \text{private care reduces mortality} \iff G_i^h > 0.$$

What does it mean to use RN as the instrument?

- R_{st} measures voucher reimbursement relative to the median private-market price; N_{st} counts private hospitals performing the surgery-year.
- The specification excludes the interaction $Z_{st} = R_{st}N_{st}$, while including the main effects of R_{st} and N_{st} as covariates.
- $Z_{st} = R_{st}N_{st}$, proxies for the effective choice set of private hospitals willing to accept publicly insured patients under the voucher scheme.
- Because Z_{st} is deterministic given the pair (R_{st}, N_{st}) , there is no nonparametric variation in Z conditional on the exact values of both components:

$$\text{Var}(Z_{st} \mid R_{st}, N_{st}, W_{st}) = 0.$$

- Identification therefore comes from a functional-form separation between additive main effects and the nonadditive interaction.

What variation remains after controlling for R and N ?

A parametric first stage may be written as

$$p(r, n, w) = \Lambda[\eta(r, n, w)], \quad \eta(r, n, w) = \alpha + g_R(r) + g_N(n) + \pi rn + w'\delta.$$

- The interaction coefficient satisfies

$$\frac{\partial^2 \eta(r, n, w)}{\partial r \partial n} = \pi.$$

It measures complementarity in the latent first-stage index (this might hint interpretation).

- In a nonlinear probability model,

$$\frac{\partial^2 p}{\partial r \partial n} = \Lambda''(\eta) [g'_R(r) + \pi n] [g'_N(n) + \pi r] + \Lambda'(\eta)\pi,$$

- Under the implemented linear-logit specification,

$$p(r, n, x) [1 - p(r, n, x)] \{ \beta_{RN} + [1 - 2p(r, n, x)] (\beta_R + \beta_{RN}n) (\beta_N + \beta_{RN}r) \}.$$

Comment. The could help state explicitly that the MTE interpretation relies on the maintained separation between additive main effects and the interaction.

From the Roy model to the latent choice index

- Let C_i denote consumption or convenience value unrelated to mortality. The conceptual Roy model is

$$V_i \approx OOP_i^0(\theta_i) - B(\theta_i) - C_i, \quad D_i = 1\{V_i \leq \nu(R_i, N_i, W_i)\}.$$

- Where θ_i indexes underlying mortality risk.
- Define the net-benefit index

$$S_i \equiv \nu(R_i, N_i, W_i) - V_i.$$

This representation is interpretive: it clarifies how the reduced-form complementarity may operate through effective financial exposure or access.

Sorting in the conceptual Roy model

- A patient chooses private care when

$$OOP^0(\theta_i) < B(\theta_i) + C_i.$$

- $B(\theta_i)$ captures willingness to pay for survival-relevant gains; C_i captures mortality-irrelevant consumption or convenience value.
- Under incomplete insurance, treatment intensity may make $OOP^0(\theta)$ increase with clinical risk.
- Sorting therefore depends on how benefits and financial costs move with θ .

When does cost exposure target high-gain patients?

Holding C fixed and taking derivatives respect to θ , the Roy index satisfies

$$S_{\theta}(\theta, C) = B_{\theta}(\theta) - OOP_{\theta}^0(\theta).$$

Positive sorting

- $B_{\theta}(\theta) > OOP_{\theta}^0(\theta)$.
- Higher-severity, higher-gain patients have greater net value from private care.
- Cost exposure may target private capacity toward high-gain patients.

Negative sorting

- $B_{\theta}(\theta) < OOP_{\theta}^0(\theta)$.
- Financial exposure rises faster than the survival-relevant benefit.
- High-gain patients may be screened out.

Takeaway. The economic issue is whether effective cost exposure tracks survival gains or financial capacity. The MTE curve will be used to answer this.

From Roy choice to the MTE margin

- Let the primitive observed shifters be $Q_i = (R_i, N_i, W_i)$, and normalize latent resistance as

$$U_i = F_{V|W}(V_i | W_i).$$

- Under the maintained first-stage specification, define

$$p(r, n, w) = \Pr(D_i = 1 | R_i = r, N_i = n, W_i = w).$$

- The treatment rule is

$$D_i = 1\{U_i \leq p(R_i, N_i, W_i)\}.$$

- The interaction $R_i N_i$ is the excluded nonlinear component under the maintained functional separation between the additive terms $g_R(R_i) + g_N(N_i)$ and the interaction term.

The Roy model explains hospital choice. The MTE interpretation then requires a defensible scalar ordering of the resulting propensity score, not an independent intervention on RN conditional on (R, N) .

MTR curves: levels, gains, and resistance

- Potential mortality under sector d is

$$MTR_d^h(u, x) = \mathbb{E}[M_i^h(d) \mid U_i = u, X_i = x].$$

- The mortality MTE and survival gain are

$$MTE^h(u, x) = MTR_1^h(u, x) - MTR_0^h(u, x),$$

$$G^h(u, x) = MTR_0^h(u, x) - MTR_1^h(u, x) = -MTE^h(u, x).$$

- Therefore,

$$\partial_u G^h(u, x) > 0 \iff \partial_u MTE^h(u, x) < 0.$$

Takeaway. An upward-sloping MTR_0 describes selection on baseline mortality risk; negative sorting on gains requires $MTR_0 - MTR_1$ to increase with resistance.

Connecting resistance to latent severity (θ): A suggested route

- Define

$$m_d^h(\theta, u, x) = \Pr(M_i^h(d) = 1 \mid \theta_i = \theta, U_i = u, X_i = x),$$

so that

$$MTR_d^h(u, x) = \mathbb{E}[m_d^h(\theta_i, u, x) \mid U_i = u, X_i = x].$$

- If,

$$U_i = q(\theta_i, x), \quad q_\theta(\theta_i, x) > 0.$$

Then $\theta_i = \theta(U_i, x)$ and $MTR_d^h(u, x) = m_d^h(\theta(u, x), x)$. Hence,

$$\partial_u MTR_0^h = \partial_\theta m_0^h \partial_u \theta,$$

and

$$\partial_u G^h = (\partial_\theta m_0^h - \partial_\theta m_1^h) \partial_u \theta.$$

Takeaway. This case makes the severity interpretation transparent, but the MTE itself does not require a one-dimensional severity ranking.

So what about non monotonicity in θ then?

- Back to the conceptual model, if $B(\theta) + C - OOP^0(\theta)$ is non-monotone in θ , the same resistance level can pool different clinical types:

$$V(\theta_1, C) = V(\theta_2, C), \quad \theta_1 \neq \theta_2.$$

- The MTE remains a valid effect along the resistance margin if the IV assumptions hold.
- What weakens is the clinical interpretation: higher U need not mean higher severity.

Comment. Resistance need not rank severity without monotonicity in θ .

Identification conditions with an interaction shifter: A possible route

- **Relevance/support:** the maintained first stage generates nondegenerate variation in

$$p(R_i, N_i, W_i).$$

- **Conditional exogeneity of the primitive shifters:**

$$(R_i, N_i) \perp (M_i^h(1), M_i^h(0), U_i) \mid W_i,$$

- **Functional-form separation:** additive effects $g_R(R_i) + g_N(N_i)$ are distinguished from the nonadditive interaction contribution.
- **Latent-index monotonicity:** if $p(r_1, n_1, w) \geq p(r_2, n_2, w)$, then $D_i(r_1, n_1) \geq D_i(r_2, n_2)$ for every latent type.
- **Continuity and overlap:** the scalar propensity score has adequate interior support.

What about IV Monotonicity with a two-component shifter?

Potential treatment is indexed by the primitive pair:

$$D_i(r, n) = 1\{S_i(r, n, W_i) \geq 0\}.$$

- Consider a potential selection index

$$S_i(r, n, W_i) = a_i(W_i) + g_R(r) + g_N(n) + \rho_i(W_i)rn,$$

- For two admissible pairs, we require

$$p(r_1, n_1, w) \geq p(r_2, n_2, w) \implies D_i(r_1, n_1) \geq D_i(r_2, n_2) \quad \forall i.$$

A sufficient condition is

$$S_i(r_1, n_1, w) \geq S_i(r_2, n_2, w) \quad \forall i$$

for every pair ordered by the maintained propensity score.

What does their linear-logit implementation impose?

The estimated first-stage index is linear in r , n , and rn :

$$\eta(r, n, x) = x'\delta + \beta_R r + \beta_N n + \beta_{RN} rn, \quad p(r, n, x) = \Lambda(\eta(r, n, x)).$$

- A threshold representation consistent with this logit is

$$D_i(r, n) = 1\{\varepsilon_i \leq \eta(r, n, X_i)\},$$

where the slopes $\beta_R, \beta_N, \beta_{RN}$ are common across patients and heterogeneity enters through the scalar threshold ε_i .

- For two pairs (r_1, n_1) and (r_2, n_2) , define

$$\Delta\eta = \beta_R \Delta r + \beta_N \Delta n + \beta_{RN} \Delta(rn).$$

Because Λ is strictly increasing,

$$p(r_1, n_1, x) \geq p(r_2, n_2, x) \iff \Delta\eta \geq 0.$$

- Since $\Delta\eta$ does not vary across latent types, $\Delta\eta \geq 0$ implies

$$D_i(r_1, n_1) \geq D_i(r_2, n_2) \quad \forall i.$$

Takeaway. The common-coefficient scalar-index logit builds pairwise monotonicity into the maintained first-stage specification; it is not established by the average first stage alone.

When is the linear index monotone in each primitive shifter?

Coordinatewise monotonicity on the observed support requires

$$\frac{\partial \eta(r, n, x)}{\partial r} = \beta_R + \beta_{RN}n \geq 0, \quad \frac{\partial \eta(r, n, x)}{\partial n} = \beta_N + \beta_{RN}r \geq 0.$$

Let $r \in [\underline{r}, \bar{r}]$ and $n \in [\underline{n}, \bar{n}]$. Then the necessary and sufficient support-wide restrictions are

$$\beta_R + \eta_{RN}n^{\min} \geq 0, \quad \beta_N + \eta_{RN}r^{\min} \geq 0,$$

where

$$n^{\min} = \begin{cases} \underline{n}, & \beta_{RN} \geq 0, \\ \bar{n}, & \beta_{RN} < 0, \end{cases} \quad r^{\min} = \begin{cases} \underline{r}, & \beta_{RN} \geq 0, \\ \bar{r}, & \beta_{RN} < 0. \end{cases}$$

- If one component rises while the other falls, coordinatewise conditions do not order treatment unless we add this condition.
- A positive β_{RN} alone is insufficient: negative main effects can overturn the interaction contribution on part of the support-It is not enough to see how the propensity score moves with RN .

Comment. The paper could report whether these derivative inequalities hold over the empirical support of (R, N) . That is the direct monotonicity implication of the implemented linear-logit first stage.

What does monotonicity require in this type of specification?

$$\Delta S_i = \Delta g_R + \Delta g_N + \rho_i \Delta(RN).$$

Common-sign interaction contribution

- $\rho_i \geq 0$ means that, within the maintained index, a larger interaction term raises the private-care index for every type.
- This is a restriction on the coefficient of the excluded nonlinear component.

Monotone treatment comparison

- A pairwise treatment ordering additionally requires $\Delta S_i \geq 0$ for every type.
- A positive $\rho_i \Delta(RN)$ can be offset by the included main-effect changes $\Delta g_R + \Delta g_N$.

Why conditional exogeneity may be plausible

- Private prices are negotiated in the privately insured market rather than for the publicly insured patients in the outcome sample.
- The primitive shifters vary at the surgery-year level, while the specification controls for surgery, year, municipality, patient characteristics, and diagnosis severity.
- Because RN is a deterministic transformation of (R, N) , exogeneity cannot be stated as if the interaction varied independently conditional on both components.
- The relevant identifying condition concerns the primitive pair, for example

$$(R_{st}, N_{st}) \perp (M_i^h(1), M_i^h(0), U_i) \mid W_i,$$

together with the functional-form restriction that separates main effects from their interaction.

Sample restrictions

- The paper focuses on serious surgical conditions and voucher-eligible patients for whom public–private choice can respond to the primitive shifters (R_i, N_i) .
- It restricts attention to Santiago and to surgeries observed across sectors and income groups.
- MTE identification requires nondegenerate variation in

$$p(R_i, N_i, W_i)$$

within this population.

- Let X_i denote the covariates indexing the MTR/MTE functions, with $W_i \subseteq X_i$. The supported estimand is

$$MTE_{\mathcal{S}}^h(u, x) = \mathbb{E}[M_i^h(1) - M_i^h(0) \mid U_i = u, X_i = x, i \in \mathcal{S}] .$$

Takeaway. The analysis identifies effects for the supported, voucher-relevant margin

Estimation algorithm I: propensity score and MTR recovery

- ❶ **First stage.** Estimate private admission by logit using the interaction and the full control vector:

$$\hat{p}_i = \Lambda\left(\hat{\alpha} + \hat{\pi}R_iN_i + \hat{g}_R(R_i) + \hat{g}_N(N_i) + W_i'\hat{\delta}\right).$$

- ❷ **Distribution regressions.** For each $d \in \{0, 1\}$, estimate by maximum likelihood

$$\Gamma_d(p, x) \equiv \Pr(Y_i = 0, D_i = d \mid P_i = p, X_i = x) = \Lambda(\beta_{0d} + x'\beta_{xd} + B_d(p)),$$

where $B_d(p)$ is a quadratic B-spline with three degrees of freedom and one internal knot at the median.

- ❸ **Differentiate.** The fitted derivative is

$$\hat{\gamma}_d(p, x) = B'_d(p)\hat{\Gamma}_d(p, x) \left[1 - \hat{\Gamma}_d(p, x)\right].$$

- ❹ **Recover the causal curves.**

$$\widehat{MTR}_1 = \hat{\gamma}_1, \quad \widehat{MTR}_0 = -\hat{\gamma}_0, \quad \widehat{MTE} = \widehat{MTR}_1 - \widehat{MTR}_0.$$

Estimation algorithm II: averaging and inference

- Evaluate $\widehat{MTR}_d(\hat{p}_i, X_i)$ for each observation on the trimmed common support and impose

$$0 \leq \widehat{MTR}_d(\hat{p}_i, X_i) \leq 1.$$

- Divide the propensity-score support into ten bins $\mathcal{B}_b$. For $N_b = \sum_i 1\{\hat{p}_i \in \mathcal{B}_b\}$, report

$$\widehat{MTR}_d(\mathcal{B}_b) = \frac{1}{N_b} \sum_{i:\hat{p}_i \in \mathcal{B}_b} \widehat{MTR}_d(\hat{p}_i, X_i), \quad \widehat{MTE}(\mathcal{B}_b) = \widehat{MTR}_1(\mathcal{B}_b) - \widehat{MTR}_0(\mathcal{B}_b).$$

- Obtain ATE, ATT, and ATU by integrating the estimated MTE with their corresponding resistance weights:

$$\hat{\tau}_k = \int_0^1 \widehat{MTE}(u) \hat{\omega}_k(u) du, \quad k \in \{ATE, ATT, ATU\}.$$

- **Inference:** run 999 municipality-level block-bootstrap replications. In every replication, re-estimate the first stage and both outcome regressions, then reconstruct the MTR, MTE, and aggregate treatment parameters. Use the bootstrap distribution for the reported 90% confidence intervals and pointwise bands.

Repeated hospitalizations: what does LIV identify?

- The paper retains each patient's last admission for a given diagnosis, but the same patient may still appear for different diagnoses.
- Let

$$\mathcal{A} = \{(i, j) : \text{retained patient-diagnosis admissions}\}, \quad \mathbb{E}_{\mathcal{A}}[H_{IJ}] = \frac{1}{|\mathcal{A}|} \sum_{(i, j) \in \mathcal{A}} H_{ij}.$$

Thus, $\mathbb{E}_{\mathcal{A}}$ averages over retained admissions, not unique patients.

- Applying LIV to the stacked rows identifies

$$MTE_A^h(u, x) = \mathbb{E}_A \left[\tau_{IJ}^h \mid U_{IJ} = u, X_{IJ} = x \right], \quad \tau_{ij}^h = M_{ij}^h(1) - M_{ij}^h(0).$$

- This is directly an effect for admissions located at resistance u ; it is not yet a patient-type effect.

Repeated hospitalizations: patient type versus admission resistance

- Decompose latent resistance into a persistent patient component and an admission-specific component:

$$V_{ij} = v(\theta_i, X_{ij}) + \omega_{ij}, \quad U_{ij} = F_{V|X}(V_{ij} | X_{ij}).$$

- The same patient shares θ_i across admissions but may have different X_{ij} , ω_{ij} , and therefore different U_{ij} .
- Define the type-specific conditional effect

$$\Delta^h(u, x, \theta) = \mathbb{E}_A \left[\tau_{IJ}^h | U_{IJ} = u, X_{IJ} = x, \theta_I = \theta \right].$$

- By iterated expectations,

$$MTE_A^h(u, x) = \int \Delta^h(u, x, \theta) f_A(\theta | u, x) d\theta.$$

The estimated MTE averages type-specific effects using the mix of patient types represented among admissions at u .

Repeated hospitalizations: interpretation of the MTE slope

- Differentiating the admission-level MTE gives

$$\begin{aligned}\partial_u MTE_A^h(u, x) &= \int \partial_u \Delta^h(u, x, \theta) f_A(\theta | u, x) d\theta \\ &\quad + \int \Delta^h(u, x, \theta) \partial_u f_A(\theta | u, x) d\theta.\end{aligned}$$

- The first term is a **within-type admission margin**; the second is a **patient-type composition margin**.
 - We can recover a meaningful effect if we impose a one-dimensional monotone bridge between admission resistance and persistent patient type
 $U_{ij} = q(\theta_i, x), \quad q_\theta(\theta_i, x) > 0.$ then $\theta_i = q^{-1}(U_{ij}, x).$
 - Hence, conditional on $U_{ij} = u$ and $X_{ij} = x$, patient type is no longer mixed:

$$F_A(\theta | U_{ij} = u, X_{ij} = x) = \delta_{q^{-1}(u, x)}(\theta).$$

- Substituting this degenerate distribution into

$$MTE_A^h(u, x) = \int \Delta^h(u, x, \theta) dF_A(\theta | u, x)$$

gives

$$MTE_A^h(u, x) = \Delta^h(u, x, q^{-1}(u, x)).$$

Repeated hospitalizations: interpretation of the MTE slope

- Differentiating the admission-level MTE gives

$$\begin{aligned}\partial_u MTE_A^h(u, x) &= \int \partial_u \Delta^h(u, x, \theta) f_A(\theta | u, x) d\theta \\ &\quad + \int \Delta^h(u, x, \theta) \partial_u f_A(\theta | u, x) d\theta.\end{aligned}$$

- The first term is a **within-type admission margin**; the second is a **patient-type composition margin**.
 - If, additionally, $\Delta^h(u, x, \theta) = G^h(\theta, x)$, then $MTE_A^h(u, x) = G^h(q^{-1}(u, x), x)$.
 - Differentiating,

$$\frac{\partial MTE_A^h(u, x)}{\partial u} = \frac{\partial_\theta G^h(q^{-1}(u, x), x)}{q_\theta(q^{-1}(u, x), x)}.$$

Because $q_\theta > 0$,

$$\text{sign} \left(\partial_u MTE_A^h(u, x) \right) = \text{sign} \left(\partial_\theta G^h(\theta, x) \right).$$

Thus, under these restrictions, variation in the estimated MTE along u directly ranks treatment gains along persistent patient type θ_i .

How to read the MTR/MTE curves

With the economic interpretation, monotonicity, and observational unit made explicit, the empirical curves have a precise reading:

- $MTR_0(p, x)$ rising with p means high-resistance patients have higher potential mortality under public care.
- $MTR_1(p, x)$ relatively flat and low means private care provides lower mortality for high-resistance patients.
- $MTE(p, x) = MTR_1(p, x) - MTR_0(p, x)$ becoming more negative means larger mortality reductions at high resistance.
- This combination is the empirical signature of negative sorting on gains.

Takeaway. The main empirical message is “private care helps most for patients least likely to use it.”

- Negative sorting in mortality risk and gains appears more pronounced among lower-income patients.
- Their MTR_0 rises more steeply and the mortality MTE becomes more negative toward the high-resistance region.
- This pattern is consistent with a higher shadow cost of out-of-pocket exposure.

Aggregate effects at 90 days

Parameter	Estimate	90% confidence interval
ATE	-0.117	$[-0.218, 0.020]$
ATT	-0.010	$[-0.083, 0.084]$
ATU	-0.173	$[-0.281, -0.010]$

- The ATT is near zero, while the ATU is substantially negative and statistically different from zero at the 10% level.
- The ordering $ATU < ATT$ is consistent with negative sorting on mortality gains.

Takeaway. Patients who remain in public care appear to have substantially larger potential mortality reductions from private admission than patients who select into it.

Mechanism: capacity or clinical ability?

- The paper studies centrally coordinated transfers when a patient's designated public hospital lacks an appropriate bed.
- Transferred patients do not choose the destination and face no out-of-pocket cost.
- Among ICU/IMCU transfers, mortality is not lower in private hospitals; the adjusted comparison points in the opposite direction.
- The comparison is suggestive

The evidence is more consistent with timely access to scarce capacity than with a uniform private-sector advantage in life-saving clinical ability.

- The results do not imply shifting all patients to private hospitals.
- They suggest relaxing financial barriers where high-gain, financially constrained patients cannot reach timely private capacity.
- In Roy terms, policy should reduce the negative association between expected gains and access to private care.

Takeaway. The relevant policy margin is targeted financial protection, not indiscriminate expansion of private utilization.

The three discussion points, restated as questions

- 1 **Economic interpretation:** Under what conditions can we interpret the MTE resistance as a interpretation of severity?
- 2 **Identification and monotonicity:** what functional-form and pairwise ordering restrictions justify a scalar MTE interpretation when RN is deterministic given (R, N) ?
- 3 **Unit of sorting:** how many patients contribute retained admissions for multiple diagnoses, and what assumptions make the admission-level resistance U_{ij} a valid monotone proxy for persistent patient type θ_i ?

Cost exposure: targeting device or barrier to high-value care?

- Cost exposure is allocatively useful only if it screens on survival gains rather than ability to pay.
- The evidence suggests that incomplete insurance may screen out high-gain patients.
- The Roy/MTE framework makes that allocation problem visible.