Online Supplementary Material for CLARA: Factual–Counterfactual Duality and Reoptimization for Linear Programming

Yoonsik Jung[a,*]

[a] *Department of Industrial and Management Engineering, Korea University, Seoul, Republic of Korea*

This document contains the full proofs of Theorems 1 and 2, Propositions 4 and 5, and technical extensions of the factual–counterfactual bound. All notation follows the main manuscript.

1. Proofs

Proof of Theorem 1 (First-order attribution)

Under the assumption that the optimal basis B is preserved, the new optimal value is $z' = (c_B + \Delta c_B)^\top B^{-1}(b + \Delta b)$. Expanding and using $z = c_B^\top B^{-1} b$, $y = c_B^\top B^{-1}$, and $x_B = B^{-1} b$ yields

$$\Delta z = y^\top \Delta b + \Delta c^\top x^* + \Delta c_B^\top B^{-1} \Delta b. \qquad \square$$

Proof of Theorem 2 (Factual–counterfactual bound)

Since $\delta^{CE} \notin \mathrm{int}(\mathcal{S}_{\mathrm{raw}})$, there exists an index i such that $-(B^{-1})_{i,:}\, \delta^{CE} \geq (x_B)_i$. By the Cauchy–Schwarz inequality,

$$(x_B)_i \leq \|(B^{-1})_{i,:}\|_2 \cdot \|\delta^{CE}\|_2,$$

so $\|\delta^{CE}\|_2 \geq (x_B)_i / \|(B^{-1})_{i,:}\|_2 \geq d_0$. $\square$

*Corresponding author

Email address: ys_jung@korea.ac.kr (Yoonsik Jung)

Proof of Proposition 4 (Cheapest basis flip)

By construction, the i^*-th primal-feasibility constraint is tight: $-(B^{-1})_{i^*,:}\,\delta^* = (x_B)_{i^*}$. For every other primal-feasibility constraint j, $\|(B^{-1})_{j,:}\|_2 \cdot d_0 \leq (x_B)_j$ by the definition of d_0, so $-(B^{-1})_{j,:}\,\delta^* \leq \|(B^{-1})_{j,:}\|_2 \cdot d_0 \leq (x_B)_j$. Thus $\delta^* \in \mathcal{S}_{\mathrm{raw}}$ and lies on its boundary. □

Proof of Proposition 5 (Chebyshev–CE chain)

The ball of radius $\min(d_0, \alpha_{\min})$ centered at the origin is contained in $\mathcal{S}$: primal-feasibility constraints are satisfied by the definition of d_0, and OAT box constraints are satisfied by the definition of $\alpha_{\min}$. Since r^* is the maximum inscribed ball radius, $r^* \geq \min(d_0, \alpha_{\min})$. For the right inequality, the ball $B(\delta^c, r^*)$ is contained in $\mathcal{S}$, and since $\mathcal{S} \subseteq \mathcal{S}_{\mathrm{raw}}$, a basis-changing perturbation $\delta^{CE} \notin \mathcal{S}_{\mathrm{raw}}$ also satisfies $\delta^{CE} \notin \mathcal{S}$, so $\|\delta^{CE} - \delta^c\|_2 \geq r^*$. By the triangle inequality, $\|\delta^{CE}\|_2 \geq r^* - \|\delta^c\|_2$. □

Proof of Proposition 5 (Chebyshev–CE chain)

The Chebyshev–CE chain referenced in the main text is given by the following proposition.

Proposition 1 (Chebyshev–CE chain)**.** *Let r^* be the Chebyshev radius of the bounded polyhedron $\mathcal{S}$ with center δ^c, and let $d^{CE} = \|\delta^{CE}\|_2$ denote the cost of a basis-changing counterfactual explanation. Then*

$$\min(d_0,\, \alpha_{\min}) \leq r^* \leq d^{CE} + \|\delta^c\|_2,$$

where $\alpha_{\min} = \min_k \min(\alpha_k^+, \alpha_k^-)$ is the smallest OAT tolerance. When the OAT bounds are non-binding for the inscribed ball ($d_0 \leq \alpha_{\min}$), the left inequality reduces to $d_0 \leq r^*$.

Proof. The ball of radius $\min(d_0, \alpha_{\min})$ centered at the origin is contained in $\mathcal{S}$: primal-feasibility constraints are satisfied by the definition of d_0, and OAT box constraints are satisfied by the definition of $\alpha_{\min}$. Since r^* is the maximum inscribed ball radius, $r^* \geq \min(d_0, \alpha_{\min})$. For the right inequality, the ball $B(\delta^c, r^*)$ is contained in $\mathcal{S}$, and since $\mathcal{S} \subseteq \mathcal{S}_{\mathrm{raw}}$, a basis-changing perturbation $\delta^{CE} \notin \mathcal{S}_{\text{raw}}$ also satisfies $\delta^{CE} \notin \mathcal{S}$, so $\|\delta^{CE} - \delta^c\|_2 \geq r^*$. By the triangle inequality, $\|\delta^{CE}\|_2 \geq r^* - \|\delta^c\|_2$. □ □

In our experiments, the OAT bounds are non-binding for every instance, so the left inequality always reduces to $d_0 \leq r^*$, justifying the use of d_0 as the operative factual–counterfactual lower bound throughout the main text.

2. Extensions of the factual–counterfactual bound

Extension to objective perturbations

For objective coefficient perturbations $\delta_c \in \mathbb{R}^n$, the basis is preserved if all reduced costs remain non-negative: $\bar{c}_j + \delta_{c_j} - (c_B + \delta_{c_B})^\top B^{-1} a_j \geq 0$ for all nonbasic j. The *dual basis robustness* is

$$d_0^{\text{dual}} = \min_{j \in \mathcal{N}} \frac{\bar{c}_j}{\|e_j - B^{-1}a_j\|_2}, \tag{1}$$

where $\mathcal{N}$ denotes the set of nonbasic indices and e_j is the j-th standard basis vector restricted to the nonbasic components of δ_c. An analogous bound holds: $d_0^{\text{dual}} \leq \|\delta_c^{CE}\|_2$ for any objective perturbation that induces a basis change. For joint (δ_b, δ_c) perturbations, the combined basis robustness is $\min(d_0, d_0^{\text{dual}})$.

Extension to alternative norms

Theorem 2 generalizes to any ℓ_p norm by replacing the Euclidean norm with the dual norm $\|\cdot\|_q$ (where $1/p + 1/q = 1$) in the denominator of the definition of d_0: $d_0^{(p)} = \min_i (x_B)_i / \|(B^{-1})_{i,:}\|_q$. In particular, for the ℓ_∞ norm (maximum absolute perturbation), $d_0^{(\infty)} = \min_i (x_B)_i / \|(B^{-1})_{i,:}\|_1$, and for the ℓ_1 norm, $d_0^{(1)} = \min_i (x_B)_i / \|(B^{-1})_{i,:}\|_\infty$. This allows the practitioner to choose the norm that best matches the perturbation model relevant to their application.

3. Scaling and normalization for simultaneous sensitivity

The absolute values of r^* and ρ are sensitive to the scaling of constraints and variables: row or column scaling changes the geometry of $\mathcal{S}$ without altering the LP solution. CLARA does not apply automatic scaling; for models with heterogeneous units we recommend one of two normalization strategies before computing the Chebyshev center.

Row-normalize.. Scale each row of A and b_i so that $\|a_i\|_2 = 1$, making r^* interpretable as the maximum uniform perturbation magnitude across normalized constraints.

Percentage-normalize.. Express perturbations as fractions of the current parameter value (δ_k/b_k), so that r^* measures the largest simultaneous percentage change.

Both strategies render ρ scale-free and comparable across models.

4. Degeneracy diagnostics

Under degeneracy, shadow prices and sensitivity ranges are basis-dependent and may not be unique (Koltai and Terlaky, 2000; Koltai and Tatay, 2011). CLARA reports three diagnostics computed from the SolveState:

- the number of degenerate basic variables, i.e. basic indices i with $|x_{B_i}| \leq \epsilon$ for a user-configurable tolerance ϵ (default 10^{-8});
- the basis condition number $\kappa(B) = \|B\|_2 \cdot \|B^{-1}\|_2$, computed from the retained factorization; and
- a per-constraint degeneracy flag identifying constraints whose shadow price is potentially non-unique (i.e. active with a zero basic variable).

When $\kappa(B)$ exceeds a user-configurable threshold (default 10^8), the Region Analyzer's output includes a cautionary note indicating that the reported shadow prices and sensitivity ranges should be interpreted with care. Enumerating all alternative optimal bases to produce tight ranges for shadow prices under degeneracy is computationally expensive and is left to future work.

5. Warm-start speedup distribution

Figure 1 shows the distribution of warm-start speedups for the 362 successfully warm-started pairs referenced in Section 6.5 of the main paper, both as pivot-count reduction (median 22.0×) and wall-clock speedup (median 3.7×). The gap between the two metrics reflects the per-pivot overhead of CLARA's educational simplex solver; for production solvers the wall-clock speedup would be closer to the pivot ratio.

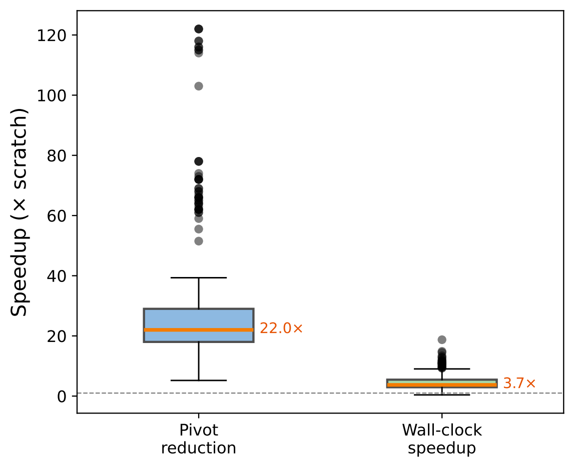

Figure 1: Distribution of warm-start speedups across the 362 matched pairs. Left: pivot-count reduction (scratch pivots / warm-start pivots), median 22.0×. Right: wall-clock speedup, median 3.7×.

6. Comparison with HiGHS warm-start

To contextualize CLARA's warm-start performance, we compared it against HiGHS's advanced-basis feature on 180 small-instance perturbation pairs. HiGHS solves the perturbed problems in a mean of 0.73 ms (cold start), while CLARA's educational solver requires 17 ms, roughly 23× slower in absolute terms. When warm-started from the old basis, HiGHS achieves only a modest 1.1× speedup (0.66 ms), because its production simplex already handles small problems efficiently and the warm-start overhead is comparable to the solve time. In contrast, CLARA achieves a 44.5× pivot reduction

Table 1: Albici reoptimization scenarios (golden test). All optimal values match the ground truth.

Scenario	Type	Method	Pivots	Optimal value	Match
base → b1	R	warm_start	0	3677.78	Yes
base → b2	R	dual simplex	1	4300.00	Yes
base → cost	C	warm_start	1	8266.67	Yes
base → columns	V	scratch	5	3744.44	Yes
compound	RC	parametric	2	10000.00	Yes

(mean 0.9 vs. 38.1 pivots) on the 62 pairs where warm-start is applicable, translating to a 4.6× wall-clock speedup within CLARA's own solver. The value of CLARA's reoptimization pipeline lies not in competing with production solvers on raw speed—for which HiGHS is clearly superior—but in the integrated explanation, attribution, and diff reports that accompany each reoptimization, none of which HiGHS provides.

7. Albici golden test

The five Albici reoptimization scenarios (Table 1) cover all change types—R (warm-start and dual), C (warm-start), V (scratch), and RC (parametric). All optimal values match the known ground truth. The pivot counts illustrate the efficiency of warm-starting: the RHS-only scenario requires 0 pivots (the old basis is already optimal for the new RHS), the dual simplex and primal warm-start each require 1 pivot, and the compound scenario requires 2 parametric breakpoints—compared to 5 pivots when the same compound problem is solved from scratch, a 60% reduction. The scratch solve for the Type V scenario (new variables) requires 5 pivots because the old basis cannot be reused when the variable set changes. This provides end-to-end validation of the entire reoptimization pipeline against an independently verified test case.

References

Koltai, T., Tatay, V., 2011. A practical approach to sensitivity analysis in linear programming under degeneracy for management decision making. International Journal of Production Economics 131, 392–398.

Koltai, T., Terlaky, T., 2000. The difference between the managerial and mathematical interpretation of sensitivity analysis results in linear programming. International Journal of Production Economics 65, 257–274.