Why Fairness Objectives Collapse on the Budget Simplex

Fairness is not chosen by the objective; it is forced by the geometry.

Meryem Essaidi

Abstract

A planner distributing a fixed budget across a population (modeled as baseline incomes) faces a familiar dilemma: should the allocation raise the worst-off agent, reduce poverty, or minimize inequality? These objectives address different aspects of disadvantage: MaxMin looks at the floor, poverty indices at the lower tail, and Gini at the whole distribution. In general, they prescribe different allocations.

We show that in the canonical direct-transfer model, these objectives agree for geometric rather than normative reasons. The feasible set of final incomes is a translated budget simplex with a least-majorized point: the final-income vector induced by ***water-filling***, which raises the poorest agents together until budget is exhausted. By the Hardy-Littlewood-Pólya theorem [6, 7], this minimizes every Schur-convex inequality objective and maximizes every Schur-concave welfare objective. As such, MaxMin maximization, Gini minimization, and FGT[4] poverty minimization agree at water-filling, as do standard inequality and welfare objectives in the same Schur classes.

This ***objective collapse*** comes from simplex geometry, not from a property of the objectives. On the budget simplex, the geometry of the feasible set resolves what would otherwise be a choice among normative criteria. More distributional information does not change the optimal allocation: ***water-filling is objective-robust.***

Style. *This note is written to make the main idea clear and easy to follow and teach. In this simple direct-transfer model, several standard fairness objectives lead to the same allocation. We explain the result from a few angles because each explanation makes a different part of the agreement easier to see. The objectives are not the same, and they do not measure fairness in the same way. The point is that, here, those differences do not change the allocation they eventually choose.*

1 The Planner's Dilemma

Consider n agents with baseline incomes $\boldsymbol{v} \in \mathbb{R}_+^n$. A planner has budget $B > 0$, chooses transfers $\boldsymbol{x} \in \mathbb{R}_+^n$, and evaluates final incomes $\boldsymbol{y} = \boldsymbol{v} + \boldsymbol{x}$. The constraint $\mathbf{1}^\top \boldsymbol{x} \leq B$ determines feasibility, not optimality. To choose an allocation, the planner must choose a criterion: e.g., raise the floor (MaxMin), reduce shortfalls (FGT), or minimize dispersion (Gini). These criteria are seldom interchangeable and can rank allocations differently. Choosing among them is a central problem in inequality measurement [1, 12, 8]. The dilemma: which fairness objective should decide the allocation?

In the direct-transfer model, this apparent dilemma collapses. The collapse is a geometric fact about the feasible set. First, for any candidate objective, some optimum exhausts the budget: unused money can be allocated uniformly, raising the floor, lowering shortfalls, and reducing relative dispersion. Second, objectives evaluate final incomes, not transfers. We thus restrict without loss to the exact-budget transfer simplex and pass from transfers to final incomes:

$$
\begin{aligned}
& \Delta_{\leq B} := \{\boldsymbol{x} \in \mathbb{R}_+^n : \mathbf{1}^\top \boldsymbol{x} \leq B\} && \text{(weak-budget transfers)} \\
\xRightarrow{\text{full-budget reduction}} \quad & \Delta_B := \{\boldsymbol{x} \in \mathbb{R}_+^n : \mathbf{1}^\top \boldsymbol{x} = B\} && \text{(exact-budget transfers)} \\
\xRightarrow{\text{translate by } \boldsymbol{v}} \quad & \mathcal{F}(\boldsymbol{v}, B) := \boldsymbol{v} + \Delta_B && \text{(translated simplex)}
\end{aligned}
$$

The geometry of $\mathcal{F}(\boldsymbol{v}, B)$ singles out a canonical point: water-filling. It raises every agent below the water level t^* to t^* and leaves every agent already above t^* unchanged, inducing $\boldsymbol{y}^*$. This point is not merely natural; it is forced: any vector in $\mathcal{F}(\boldsymbol{v}, B)$ has the same total, so feasible vectors differ only in how this total is distributed, and majorization orders apply: $\boldsymbol{y}^* \preceq \boldsymbol{y}$ means $\boldsymbol{y}^*$ is weakly more equal than $\boldsymbol{y}$, i.e., *obtainable from $\boldsymbol{y}$ by a finite sequence of Pigou-Dalton transfers.* Since $\boldsymbol{y} \neq \boldsymbol{y}^*$, some coordinate lies below t^* and some compensating coordinate lies above it. Moving income from the latter to the former is a progressive transfer that preserves feasibility and moves $\boldsymbol{y}$ closer to $\boldsymbol{y}^*$. Iterating gives $\boldsymbol{y}^*$, thus every other feasible final-income vector is weakly more unequal than water-filling.

$$
\begin{aligned}
& \mathcal{F}(\boldsymbol{v}, B) := \boldsymbol{v} + \Delta_B, \ \mathbf{1}^\top \boldsymbol{y} = \mathbf{1}^\top \boldsymbol{v} + B, && \text{(simplex = fixed-sum geometry)} \\
\xRightarrow{\text{majorization}} \quad & \preceq := \text{the equality order on fixed-sum vectors} && \text{(compare feasible incomes)} \\
\xRightarrow{\text{water-filling}} \quad & \boldsymbol{y}^* \preceq \boldsymbol{y} \qquad \forall \boldsymbol{y} \in \mathcal{F}(\boldsymbol{v}, B) && \text{(most equal feasible point).}
\end{aligned}
$$

The objective collapse follows from the Schur–majorization correspondence. Schur-convex objectives decrease and Schur-concave objectives increase under equalizing transfers. Since $\boldsymbol{y}^*$ is the unique least-majorized point of $\mathcal{F}(\boldsymbol{v}, B)$, water-filling minimizes every Schur-convex objective and maximizes every Schur-concave objective on this set.

$$
\begin{aligned}
\boldsymbol{y}^* \preceq \boldsymbol{y} \quad \forall \boldsymbol{y} \in \mathcal{F}(\boldsymbol{v}, B) \quad & \xRightarrow{\text{Schur order}} \begin{cases} \phi(\boldsymbol{y}^*) \leq \phi(\boldsymbol{y}), & \phi \text{ Schur-convex}, \\ \psi(\boldsymbol{y}^*) \geq \psi(\boldsymbol{y}), & \psi \text{ Schur-concave}, \end{cases} && \text{(objective inequalities)} \\
& \xRightarrow{\text{Optimization}} \begin{cases} \boldsymbol{y}^* \in \operatorname{argmin}_{\boldsymbol{y} \in \mathcal{F}(\boldsymbol{v}, B)} \phi, \\ \boldsymbol{y}^* \in \operatorname{argmax}_{\boldsymbol{y} \in \mathcal{F}(\boldsymbol{v}, B)} \psi, \end{cases} && \text{(universal optimality).}
\end{aligned}
$$

By the Hardy–Littlewood–Pólya theorem, this geometric fact becomes an objective collapse: MaxMin maximization, Gini minimization, and $\mathrm{FGT}_{\alpha,r}$ poverty minimization for $\alpha \geq 1$ all select the same transfer:

$$\boxed{\text{On the direct-transfer budget simplex, changing the fairness objective does not change the optimal allocation.}}$$

Contribution. *The majorization theorem is classical [6, 7, 9]. The contribution of this note is the fair-allocation translation: on the direct-transfer budget simplex, a choice among fairness objectives becomes a geometric statement about the least-majorized feasible final-income vector. MaxMin, Gini, and poverty minimization express different fairness principles. They agree here because the feasible set has already selected water-filling.*

2 Direct Transfers and the Budget Simplex

We formalize the first reduction: a planner starts from baseline $\boldsymbol{v} \in \mathbb{R}^n_+$, chooses transfers $\boldsymbol{x} \in \mathbb{R}^n_+$, and evaluates the induced outcome $\boldsymbol{y} = \boldsymbol{v} + \boldsymbol{x}$. The constraint $\mathbf{1}^\top \boldsymbol{x} \le B$ is weak: any residual budget can be allocated uniformly or given to the poorest, thus raising the floor, lowering shortfalls, or reducing dispersion. We thus reduce the set to its exact-budget face. As objectives evaluate final incomes rather than transfers, this simplex is then translated by $\boldsymbol{v}$.

Lemma 2.1 (Full-budget reduction)**.** *For MaxMin, Gini, and* $\mathrm{FGT}_{\alpha,r}$ *poverty optimization with* $\alpha > 1$*, there exists an optimal transfer that exhausts the budget. Equivalently, each of these optimization problems admits an optimum in* Δ_B*.*

Proof. Let $\boldsymbol{x}$ satisfy $\mathbf{1}^\top \boldsymbol{x} < B$, and set $\varepsilon = B - \mathbf{1}^\top \boldsymbol{x}$. Start from $\boldsymbol{v} + \boldsymbol{x}$ and allocate the ε by water-filling the current min-income agents. This raises the min, so MaxMin improves. It also moves the final-income vector to equality via *progressive transfers*, so Schur-convex inequality measures (e.g. Gini) decrease. FGT poverty decreases too as the allocation targets agents with positive shortfall until the budget is exhausted or poverty is eliminated. □

$$\begin{aligned} &\Delta_{\le B} := \{\boldsymbol{x} \in \mathbb{R}^n_+ : \mathbf{1}^\top \boldsymbol{x} \le B\} \qquad \text{(weak-budget transfers)} \\ \xrightarrow{\text{full-budget reduction}}\; &\Delta_B := \{\boldsymbol{x} \in \mathbb{R}^n_+ : \mathbf{1}^\top \boldsymbol{x} = B\} \qquad \text{(exact-budget transfers)} \end{aligned}$$

Remark 2.2 (Final incomes, not transfers). Comparisons are made on $\boldsymbol{y} = \boldsymbol{v} + \boldsymbol{x}$, not $\boldsymbol{x}$. Water-filling may look unequal as a transfer vector as poorer agents receive larger transfers. The induced final-income vector is what matters. Vectors in $\mathcal{F}(\boldsymbol{v}, B)$ have the same total, and differ only in how this total is distributed. This is the key simplification and natural domain for majorization: *one vector is more equal than another if it can be obtained by moving income from higher to lower coordinates.* The problem then becomes geometric: among all final-incomes, find the most equal one.

$$\begin{aligned} \xrightarrow{\text{translate by } \boldsymbol{v}}\; &\mathcal{F}(\boldsymbol{v}, B) := \boldsymbol{v} + \Delta_B = \{\boldsymbol{v} + \boldsymbol{x} : \boldsymbol{x} \in \Delta_B\} \qquad \text{(feasible final incomes)} \\ \Longrightarrow\; &\mathbf{1}^\top \boldsymbol{y} = \mathbf{1}^\top \boldsymbol{v} + B \qquad (\boldsymbol{y} \in \mathcal{F}(\boldsymbol{v}, B)). \qquad \text{(fixed-sum geometry)} \end{aligned}$$

The canonical candidate is water-filling. The water level t^* is the unique level satisfying $\sum_i (t^* - v_i)^+ = B$. The transfer $\boldsymbol{x}^*$ raises agents below t^* up to t^*, leaves agents already above t^* unchanged, and exhausts the budget. Formally: for $(\boldsymbol{v}, B)$, let t^*, $\boldsymbol{x}^*$, and $\boldsymbol{y}^*$ denote the threshold, transfer vector, and final-income vector given by

$$\sum_i (t^* - v_i)^+ = B, \qquad x_i^* = (t^* - v_i)^+ \ \forall i \in [n], \qquad \boldsymbol{x}^* = (t^* \mathbf{1} - \boldsymbol{v})_+, \qquad \boldsymbol{y}^* = \boldsymbol{v} + \boldsymbol{x}^*, \qquad y_i^* = \max\{v_i, t^*\}. \tag{1}$$

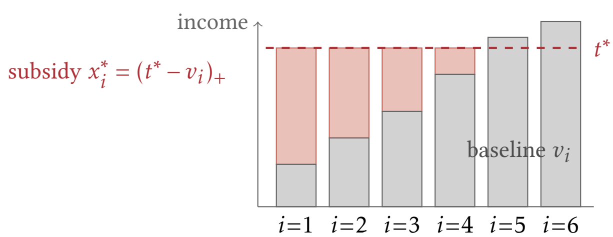

Figure 1: Water-filling raises the floor to t^*. Gray bars are baseline incomes; maroon are subsidies.

The next section proves the central geometric fact: $\boldsymbol{y}^*$ is the unique least-majorized point of $\mathcal{F}(\boldsymbol{v}, B)$. Thus the reduction and translation turn the objective-choice problem into a fixed-sum equality problem on a simplex.

$$\begin{aligned} &\xrightarrow{\text{apply majorization} \boldsymbol{v}}\; \preceq := \text{equality order on the feasible set} \\ &\xrightarrow{\text{reveals candidate}}\; \text{water-filling is most-equal} \end{aligned}$$

3 Water-Filling Is Least-Majorized

The previous section reduced the planner's problem to the fixed-total final-income set $\mathcal{F}(\boldsymbol{v}, B) = \boldsymbol{v} + \Delta_B$ and defined the water-filled vector $\boldsymbol{y}^*$. We now prove a geometric fact: $\boldsymbol{y}^*$ is the unique most equal point of $\mathcal{F}(\boldsymbol{v}, B)$. For $\boldsymbol{y} \in \mathbb{R}^n$, let $y_1^{\downarrow} \geq \cdots \geq y_n^{\downarrow}$ be its coordinates in decreasing order. For equal-sum $\boldsymbol{y}, \boldsymbol{z} \in \mathbb{R}^n$, $\boldsymbol{z} \preceq \boldsymbol{y}$ if $\boldsymbol{z}$ is more equal than $\boldsymbol{y}$, i.e.

$$\boldsymbol{z} \preceq \boldsymbol{y} \iff \sum_{i=1}^{k} z_i^{\downarrow} \leq \sum_{i=1}^{k} y_i^{\downarrow} \qquad k = 1, \ldots, n-1, \tag{majorization}$$

$$\boldsymbol{z} \preceq \boldsymbol{y} \iff \boldsymbol{z} \text{ is obtained from } \boldsymbol{y} \text{ by finitely many Pigou–Dalton transfers}, \tag{HLP}$$

$$\boldsymbol{y}^* \preceq \boldsymbol{y} \quad \forall \boldsymbol{y} \in \mathcal{F}(\boldsymbol{v}, B) \iff \boldsymbol{y}^* \text{ is least-majorized on } \mathcal{F}(\boldsymbol{v}, B). \tag{geometric target}$$

A least-majorized point of a fixed-sum set is feasible and weakly more equal than every other feasible point. We use the HLP characterization [6, 7]: for equal-sum vectors, $\boldsymbol{z} \preceq \boldsymbol{y}$ *iff* $\boldsymbol{z}$ can be obtained from $\boldsymbol{y}$ by finite Pigou-Dalton transfers, each moving income from a higher to a lower coordinate without reversing order. Thus, to prove $\boldsymbol{y}^*$ is least-majorized, it suffices to show that any other vector can be converted into $\boldsymbol{y}^*$ via finitely many such moves.

Theorem 3.1 (Water-filling is least-majorized)**.** *Let $\boldsymbol{v} \in \mathbb{R}^n_+$, $B > 0$, and $\mathcal{F}(\boldsymbol{v}, B) = \boldsymbol{v} + \Delta_B$. The water-filled vector $\boldsymbol{y}^*$ is the unique least-majorized point of $\mathcal{F}(\boldsymbol{v}, B)$, and strictly so if $\boldsymbol{y} \neq \boldsymbol{y}^*$:*

$$\boldsymbol{y}^* \preceq \boldsymbol{y} \qquad \forall \boldsymbol{y} \in \mathcal{F}(\boldsymbol{v}, B).$$

Proof. Fix $\boldsymbol{y} \in \mathcal{F}(\boldsymbol{v}, B)$. If $\boldsymbol{y} = \boldsymbol{y}^*$, there is nothing to prove. Otherwise, $\boldsymbol{y}$ and $\boldsymbol{y}^*$ have the same total income, so $\boldsymbol{y}$ has both deficits and surpluses relative to $\boldsymbol{y}^*$. Define these non-empty deficit and surplus sets:

$$D = \{i : y_i < y_i^*\}, \qquad S = \{j : y_j > y_j^*\}.$$

The water level separates them: every deficit coordinate lies below t^*, and every surplus coordinate lies above t^*:

$$\left.\begin{array}{ll} i \in D \Longrightarrow y_i^* = t^* \Longrightarrow y_i < y_i^* = t^*, & \text{(deficit below water)} \\ j \in S \Longrightarrow y_j > y_j^* = \max\{v_j, t^*\} \geq t^*. & \text{(surplus above water)} \end{array}\right\} \Longrightarrow \quad y_i < t^* < y_j \quad (i \in D,\ j \in S).$$

Choose $i \in D$ and $j \in S$. Move the largest amount from j to i without crossing: $\delta = \min\{y_i^* - y_i,\ y_j - y_j^*\} > 0$, $\widetilde{\boldsymbol{y}} = \boldsymbol{y} + \delta e_i - \delta e_j$. This is an equalizing move that preserves total income and feasibility, while fixing at least one coordinate:

$$\left.\begin{array}{ll} \widetilde{y}_i = y_i + \delta \leq y_i^* = t^*, \quad \widetilde{y}_j = y_j - \delta \geq y_j^* \geq t^*, & \text{(no crossing)} \\ \widetilde{y}_i = y_i^* \qquad \text{or} \qquad \widetilde{y}_j = y_j^*. & \text{(one coordinate fixed)} \end{array}\right\} \Longrightarrow \quad \widetilde{\boldsymbol{y}} \text{ is obtained from } \boldsymbol{y}.$$

Repeat, recomputing D and S after each step. Each step preserves feasibility and total income, is a Pigou–Dalton transfer, and fixes at least one coordinate. The process terminates after finitely many steps at $\boldsymbol{y}^*$. By the HLP characterization, $\boldsymbol{y}^* \preceq \boldsymbol{y}$. If $\boldsymbol{y} \neq \boldsymbol{y}^*$, the first move is nontrivial so $\boldsymbol{y}^* \prec \boldsymbol{y}$ and no other vector is least-majorized. ☐

The theorem is the geometric engine. On the translated budget simplex, water-filling is not merely a natural allocation rule. It induces the unique feasible final-income vector that is weakly more equal than every other feasible final-income vector. The next section turns this least-majorization fact into objective agreement.

4 Objective Collapse

Section 3 identified the geometric primitive: $\boldsymbol{y}^*$ is the unique least-majorized point of $\mathcal{F}(\boldsymbol{v}, B)$. We now translate this into optimization. This is the Schur-majorization correspondence: objectives that decrease under equalizing transfers are minimized at least-majorized points, and objectives that increase under equalizing transfers are maximized there.

Definition 4.1 (Schur-convexity and Schur-concavity)**.** A function $\phi : \mathbb{R}^n \to \mathbb{R}$ is *Schur-convex* if $\boldsymbol{z} \preceq \boldsymbol{y} \implies \phi(\boldsymbol{z}) \le \phi(\boldsymbol{y})$, and *Schur-concave* if $\boldsymbol{z} \preceq \boldsymbol{y} \implies \phi(\boldsymbol{z}) \ge \phi(\boldsymbol{y})$. It is strict if $\boldsymbol{z} \prec \boldsymbol{y}$ implies strict inequality.

Schur-convex objectives formalize inequality criteria that improve under Pigou-Dalton transfers [2]: a transfer from a richer to a poorer coordinate decreases their value. Schur-concave objectives formalize welfare criteria that weakly improve under the same transfers when total income is fixed. The only fact we need is immediate from the definition:

$$\begin{aligned} &\boldsymbol{z} \preceq \boldsymbol{y} \Longrightarrow \phi(\boldsymbol{z}) \le \phi(\boldsymbol{y}), && \text{(Schur-convexity)} \\ &\boldsymbol{z} \preceq \boldsymbol{y} \Longrightarrow \psi(\boldsymbol{z}) \ge \psi(\boldsymbol{y}), && \text{(Schur-concavity)} \\ &\boldsymbol{z} \prec \boldsymbol{y} \Longrightarrow \text{strict inequality}, && \text{(strict Schur property).} \end{aligned}$$

Inequality measures (Gini, Theil, variance, generalized entropy) are Schur-convex [13, 3]; welfare measures (MaxMin, concave-sum, Atkinson)[1, 13, 12] are Schur-concave [10, 1]. FGT poverty FGT_α[1] is Schur-convex[2] [4],

Proposition 4.2 (Least-majorized points are Schur-optimal)**.** *Let $\mathcal{C} \subseteq \mathbb{R}^n$ be a fixed-total feasible set, and $\boldsymbol{y}^\circ \in \mathcal{C}$ satisfy $\boldsymbol{y}^\circ \preceq \boldsymbol{y}$, $\forall \boldsymbol{y} \in \mathcal{C}$. Then $\boldsymbol{y}^\circ$ minimizes every Schur-convex objective and maximizes every Schur-concave objective over $\mathcal{C}$. If the objective is strict wrt majorization and $\boldsymbol{y}^\circ \prec \boldsymbol{y}$ for every $\boldsymbol{y} \neq \boldsymbol{y}^\circ$, then $\boldsymbol{y}^\circ$ is the unique optimizer.*

For Schur-convex ϕ, $\boldsymbol{y}^\circ \preceq \boldsymbol{y}$ gives $\phi(\boldsymbol{y}^\circ) \le \phi(\boldsymbol{y})$, $\forall \boldsymbol{y} \in \mathcal{C}$, so $\boldsymbol{y}^\circ$ minimizes ϕ over $\mathcal{C}$. For Schur-concave ψ, the same assumption gives $\psi(\boldsymbol{y}^\circ) \ge \psi(\boldsymbol{y})$, $\forall \boldsymbol{y} \in \mathcal{C}$, so $\boldsymbol{y}^\circ$ maximizes ψ over $\mathcal{C}$. Strictness gives uniqueness. Applying this proposition to the translated simplex gives the universal optimality statement.

Theorem 4.3 (Universal optimality)**.** *Over $\mathcal{F}(\boldsymbol{v}, B) = \boldsymbol{v} + \Delta_B$, the water-filled vector $\boldsymbol{y}^*$ minimizes every Schur-convex and maximizes every Schur-concave objective. If the objective is strict wrt majorization, then $\boldsymbol{y}^*$ is the unique optimizer.*

By Theorem 3.1, $\boldsymbol{y}^* \preceq \boldsymbol{y}$, $\forall \boldsymbol{y} \in \mathcal{F}(\boldsymbol{v}, B)$, with strict majorization whenever $\boldsymbol{y} \neq \boldsymbol{y}^*$. Since $\mathcal{F}(\boldsymbol{v}, B)$ is fixed-total, Proposition 4.2 applies. We now specialize to three canonical objectives: MaxMin, Gini, and $\mathrm{FGT}_{\alpha,r}$. For $\bar{y} = \frac{1}{n}\sum_i y_i$:

$$\mathrm{MaxMin}(\boldsymbol{y}) = \min_i y_i, \qquad \mathrm{Gini}(\boldsymbol{y}) = \frac{1}{2n^2\bar{y}} \sum_{i,j} |y_i - y_j|, \qquad \mathrm{FGT}_{\alpha,r}(\boldsymbol{y}) = \frac{1}{n}\sum_i \left(\frac{(r - y_i)^+}{r}\right)^\alpha.$$

The objective collapse thus follows from the universal theorem. MaxMin, Gini, and poverty do not encode the same normative principle. They agree here because the direct-transfer budget simplex has a unique least-majorized final-income vector, and every Schur-respecting objective recognizes that same geometric point. In this model, objective choice becomes simplex geometry, and simplex geometry selects water-filling.

$$\boxed{\boldsymbol{x}^* \in \operatorname*{argmax}_{\boldsymbol{x} \in \Delta_B} \mathrm{MaxMin}(\boldsymbol{v} + \boldsymbol{x}) \cap \operatorname*{argmin}_{\boldsymbol{x} \in \Delta_B} \mathrm{Gini}(\boldsymbol{v} + \boldsymbol{x}) \cap \operatorname*{argmin}_{\boldsymbol{x} \in \Delta_B} \mathrm{FGT}_{\alpha,r}(\boldsymbol{v} + \boldsymbol{x})}$$

[1] FGT is strictly Schur-convex for $\alpha > 1$ when poverty is not fully eliminated.

[2] For FGT with threshold r, uniqueness statements refer to the nondegenerate case $\sum_i (r - v_i)^+ > B$.

5 The Geometry of Robustness

In the canonical direct-transfer model, the planner starts with an apparent objective-choice problem. Passing to final incomes reduces the feasible choices to $v + \Delta_B$. The simplex has a unique most-equal point, water-filling. Majorization makes this point optimal for every Schur-respecting objective. The named objectives agree because they are Schur-respecting, and the simplex gives them a common point to select. Hence the objective choice collapses:

$$\arg \text{Fairness}(v + \Delta_B) = \{y^*\} \qquad \text{for every strict Schur-respecting criterion.}$$

For non-strict criteria, water-filling remains optimal, though other optima may also exist. In this sense, water-filling is not a compromise among fairness objectives. It is the geometric solution to the budget-simplex problem. The objectives agree because the feasible set has already selected the most equal final-income vector.

The objective collapse doesn't mean MaxMin, FGT, and Gini encode the same criterion. They don't. These objectives disagree on many sets, and choosing among them is needed in inequality measurement. Our point is narrower: on the direct-transfer budget simplex, those normative distinctions do not translate into allocation distinctions.

This also identifies the primitive source of the result. The collapse comes from the feasible set $\boldsymbol{v} + \Delta_B$. The proof uses three structural features: transfers are direct, transfers are nonnegative, and the budget is fixed. Together, these make the feasible final-income set a translated simplex. The translated simplex has a single least-majorized point, and the HLP correspondence forces every Schur-respecting objective to optimize there.

This also explains why the collapse is conditional. If the set is perturbed by taxes, take-up frictions, household sharing, or eligibility constraints, a least-majorized point need not exist, and objectives can again select different policies.

Operations research and discrete optimization [14, 5, 11] study least-majorized points. We use that geometry to explain a fair-allocation phenomenon: when the set has a majorization-minimal point, objective choice doesn't change allocation. ***On the budget simplex, the feasible set has already chosen the fairest allocation.***

References

[1] Anthony B. Atkinson. On the measurement of inequality. *Journal of Economic Theory*, 2(3):244–263, 1970.

[2] Hugh Dalton. The measurement of the inequality of incomes. *The Economic Journal*, 30(119):348–361, 1920.

[3] Dasgupta, Sen, and Starrett. Notes on the measurement of inequality. *Journal of Economic Theory*, 6(2):180–187, 1973.

[4] James Foster, Joel Greer, and Erik Thorbecke. A class of decomposable measures. *Econometrica*, 52(3):761–766, 1984.

[5] A. Goel and A. Meyerson. Simultaneous optimization for concave functions. *J. Comput. Syst. Sci.*, 72(5):854–866, 2006.

[6] G. H. Hardy, J. E. Littlewood, and G. Pólya. Simple inequalities satisfied by convex functions. *Messenger of Mathematics*, 58:145–152, 1929.

[7] G. H. Hardy, J. E. Littlewood, and G. Pólya. *Inequalities.* Cambridge University Press, 2 edition, 1952.

[8] Serge-Christophe Kolm. Unequal inequalities. i. *Journal of Economic Theory*, 12(3):416–442, 1976.

[9] Albert W. Marshall, Ingram Olkin, and Barry C. Arnold. *Inequalities: Theory of Majorization.* Springer, 2 edition, 2011.

[10] John Rawls. *A Theory of Justice.* Harvard University Press, 1971.

[11] M. Schoot Uiterkamp, J. Klundert, and G. Post. Least-majorized elements of polyhedra. *Math. Oper. Res.*, 2024.

[12] Amartya Sen. *On Economic Inequality.* Oxford University Press, 1973.

[13] Anthony F. Shorrocks. The class of additively decomposable inequality measures. *Econometrica*, 48(3):613–625, 1980.

[14] Arie Tamir. The least element property of hierarchical structures. *Mathematical Programming*, 68:341–349, 1995.

AI disclosure. The authors used Claude for copyediting, LaTeX formatting, and organization of author-written material. All mathematical claims, proofs, examples, citations, and final text were developed, verified, and finalized by the authors, who take full responsibility for the paper.

A Four-step summary

Step 1

We study n agents with baseline incomes $\boldsymbol{v} \in \mathbb{R}_+^n$. A planner has a fixed budget $B > 0$ and may add nonnegative transfers to these incomes. A transfer vector $\boldsymbol{x} \in \mathbb{R}_+^n$ induces final incomes $\boldsymbol{y} = \boldsymbol{v} + \boldsymbol{x}$, and the planner's objectives are evaluated on $\boldsymbol{y}$, not on $\boldsymbol{x}$. The budget constraint $\mathbf{1}^\top \boldsymbol{x} \le B$ determines which transfers are feasible but does not select an allocation.

To select among feasible transfers, the planner must choose a fairness criterion. MaxMin targets the poorest agent, poverty minimization targets shortfalls below a threshold, and Gini minimization targets global dispersion. These criteria measure different features of the final-income distribution, and, in general, can select different allocations. Thus the planner appears to face an objective-choice problem: which fairness criterion should determine the allocation?

The main result shows that, in the direct-transfer model, this choice collapses. The collapse does not come from moral equivalence among the objectives. It comes from the geometry of the feasible set. For the objectives considered here, some optimum exhausts the budget; exact-budget transfers form the simplex Δ_B; and final incomes form the translated simplex $\boldsymbol{v} + \Delta_B$. On this simplex, water-filling is the unique least-majorized point, and every Schur-respecting objective selects it.

Step 2

The first step is to remove an inessential degree of freedom. The planner's constraint is $\mathbf{1}^\top \boldsymbol{x} \le B$, but unused budget is never helpful for most objectives. Starting from any feasible final-income vector, leftover budget can be spent by either raising the poorest or allocating uniformly. This raises the floor, lowers shortfalls, or reduces dispersion. In general, any monotone objective admits a full-budget optimum and we thus restrict to the exact-budget simplex:

$$\Delta_B = \{\boldsymbol{x} \in \mathbb{R}_+^n : \mathbf{1}^\top \boldsymbol{x} = B\}.$$

Since the objectives are evaluated on final incomes, the relevant set is not Δ_B itself but its translation by baseline incomes:

$$\mathcal{F}(\boldsymbol{v}, B) = \boldsymbol{v} + \Delta_B.$$

Every vector in $\mathcal{F}(\boldsymbol{v}, B)$ has total income $\mathbf{1}^\top \boldsymbol{v} + B$, so feasible final-income vectors differ only in how this total is distributed.

Step 3

The fixed-sum geometry of equal-sum vectors allows majorization comparisons, which order distributions by inequality: $\boldsymbol{z} \preceq \boldsymbol{y}$ means that $\boldsymbol{z}$ is more equal than $\boldsymbol{y}$, i.e., that $\boldsymbol{z}$ can be obtained from $\boldsymbol{y}$ by finite progressive transfers. Thus, the planner's problem becomes geometric: find the most equal point in $\mathcal{F}(\boldsymbol{v}, B)$. *Water-filling is the natural candidate.*

The main theorem shows that this intuitive allocation is forced by the simplex geometry: the most equalizing rule raises lowest incomes together to a common level t^* until budget is exhausted; any other feasible $\boldsymbol{y} \in \mathcal{F}(\boldsymbol{v}, B)$ majorizes it: $\boldsymbol{y}^* \preceq \boldsymbol{y}$. Formally, the water level, induced transfer, and final-income vectors are:

$$\sum_i (t^* - v_i)^+ = B, \qquad x_i^* = (t^* - v_i)^+, \qquad y_i^* = v_i + x_i^* = \max\{v_i, t^*\}.$$

The proof. If $\boldsymbol{y} \neq \boldsymbol{y}^*$, then some are below t^* and some above. Moving mass from a surplus to a deficit coordinate is a feasible progressive transfer that moves $\boldsymbol{y}$ closer to $\boldsymbol{y}^*$. Iterating reaches $\boldsymbol{y}^*$, so $\boldsymbol{y}$ is more unequal than $\boldsymbol{y}^*$, i.e. $\boldsymbol{y}^*$ is least-majorized.

Step 4

The objective collapse follows from the Schur-majorization correspondence. Schur-convex (resp. Schur concave) objectives decrease (resp. increase) under equalizing transfers. Since water-filling $\boldsymbol{x}^*$ induces the least-majorized outcome, it optimizes every Schur-respecting objective; with MaxMin, Gini, and $\mathrm{FGT}_{\alpha,r}$ as canonical examples. The conclusion is allocation-equivalence, not normative equivalence. MaxMin, Gini, and poverty minimization do not express the same fairness principle. They agree here because the direct-transfer budget simplex has a unique least-majorized final-income vector. In this model, objective choice becomes simplex geometry, and simplex geometry selects water-filling.

B Step by Step breakdown

A planner has a limited budget to improve unequal baseline incomes, must select an allocation, and evaluate the induced outcome.

Setup

1. **Start with direct transfers.**
   We study n agents with baseline incomes $\boldsymbol{v} \in \mathbb{R}^n_+$. A planner has budget $B > 0$ and chooses nonnegative transfers $\boldsymbol{x} \in \mathbb{R}^n_+$. Transfers induce final incomes: each transfer vector $\boldsymbol{x}$ induces $\boldsymbol{y} = \boldsymbol{v} + \boldsymbol{x}$. Objectives are evaluated on final incomes, not transfers: The planner's objective depends on $\boldsymbol{y}$, not on $\boldsymbol{x}$ itself.
2. **Feasibility alone does not select an allocation.**
   The budget constraint defines which transfers are allowed, but not which one is optimal. The planner needs a criterion: MaxMin, FGT, Gini are standard candidates. These criteria can disagree: their optimizers need not coincide. This creates a dilemma.

Apparent problem is objective choice: Which fairness criterion should determine the allocation?

Prune

3. **First reduction: spend the whole budget.**
   For the named monotone objectives, unused budget cannot improve the objective. There exists an optimum with $\mathbf{1}^\top \boldsymbol{x} = B$. Hence we may restrict to $\Delta_B = \{\boldsymbol{x} \in \mathbb{R}^n_+ : \mathbf{1}^\top \boldsymbol{x} = B\}$.
4. **Second reduction: look at final incomes directly.**
   Since objectives evaluate $\boldsymbol{y} = \boldsymbol{v} + \boldsymbol{x}$, we analyze the final-income simplex $\mathcal{F}(\boldsymbol{v}, B) = \boldsymbol{v} + \Delta_B = \{\boldsymbol{v} + \boldsymbol{x} : \boldsymbol{x} \in \Delta_B\}$. Now all feasible final-income vectors have the same total. Fixed totals make equality comparable: equal-sum vectors can be ordered by majorization: $\boldsymbol{z} \preceq \boldsymbol{y}$ means $\boldsymbol{z}$ is weakly more equal than $\boldsymbol{y}$.

Simplex Geometry allows majorization comparisons which reveal a candidate solution: Water-filling.

Optimize

5. **Water-filling is the candidate most-equal point.**
   Water-filling raises the lowest incomes together until the budget is exhausted. The water level t^* satisfies $\sum_i (t^* - v_i)^+ = B$. It raises everyone below the water level and leaves everyone above it unchanged: $x_i^* = (t^* - v_i)^+,\ y_i^* = v_i + x_i^* = \max\{v_i, t^*\}$.
6. **Main theorem: water-filling is forced by simplex geometry.**
   For every feasible final-income vector $\boldsymbol{y} \in \mathcal{F}(\boldsymbol{v}, B)$, $\boldsymbol{y}^* \preceq \boldsymbol{y}$. Thus water-filling is the unique least-majorized point of $\mathcal{F}(\boldsymbol{v}, B)$. Compare any feasible vector to water-filling. Define deficit and surplus sets: deficits lie below the water level; surpluses lie above it; i.e., if $i \in D$, then $y_i < t^*$, if $j \in S$, then $y_j > t^*$. Move mass from surplus to deficit. This is a Pigou–Dalton transfer. Iterate until water-filling is reached. Each step fixes at least one coordinate, preserves feasibility and total sum, and makes the vector more equal. So finitely many steps reach $\boldsymbol{y}^*$.

Majorization establishes water-filling as most-equal: Objective collapse follows from Schur theory.

Translate

7. **Schur theory translates this to optimality.**
   Schur-convex objectives decrease under equalizing transfers, while Schur-concave objectives increase under equalizing transfers. Water-filling optimizes every Schur-respecting objective: since $\boldsymbol{y}^* \preceq \boldsymbol{y}$ for every feasible $\boldsymbol{y}$, water-filling minimizes every Schur-convex objective and maximizes every Schur-concave objective over $\mathcal{F}(\boldsymbol{v}, B)$. MaxMin, Gini, and FGT are illustrative instances.
8. **The objectives agree because the feasible set already picked the allocation.**
   The result is allocation-equivalence, not normative equivalence: Direct transfers turn objective choice into simplex geometry, and simplex geometry selects water-filling.

This choice collapses for direct transfers: Objectives select the same allocation as the feasible set has special geometry.

C TL;DR

A budget-constrained planner appears to face two constraints. The first is physical: the planner can spend only B. The second is normative: the planner must decide which fairness objective to optimize. Should the allocation raise the floor, reduce poverty, or minimize inequality? On a general feasible set, these choices can lead to different allocations. The first simplification is to reduce to full-budget allocations and look at their induced final incomes rather than the transfers themselves. A transfer x matters only through the final vector $y = v + x$. For the named objectives below, exact-budget transfers are without loss: unused budget can be spent on low-income agents or split uniformly without worsening the objective. Thus the planner's feasible final-income choices reduce to the translated budget simplex

$$v + \Delta_B = \{v + x : x \geq 0,\ \mathbf{1}^\top x = B\}.$$

This turns the allocation problem into a geometric question: which point of this simplex should the planner choose? The second simplification is to use majorization rather than comparing/evaluating objectives. Viewed geometrically, one point of $v + \Delta_B$ stands out: the water-filled vector y^*, obtained by raising the poorest incomes together until the budget is exhausted. The main theorem proves that this point is the unique least-majorized point of the simplex:

$$y^* \preceq y \qquad \text{for every } y \in v + \Delta_B.$$

Equivalently, every other feasible final-income vector can be moved toward y^* by Pigou–Dalton transfers from higher coordinates to lower coordinates. The final step is the Schur bridge. If a fixed-sum feasible set has a least-majorized point, then every Schur-convex objective is minimized there and every Schur-concave objective is maximized there. Thus the objective choice collapses on $v + \Delta_B$: MaxMin, Gini minimization, and $\mathrm{FGT}_{\alpha,r}$ minimization for $\alpha \geq 1$ all select water-filling, subject to the usual strictness and nondegeneracy qualifications. The reason is not that these objectives are morally equivalent. The reason is that the direct-transfer budget simplex has already selected a unique most-equal final-income vector.

$$\begin{array}{rcl}
\text{objective choice} & \rightsquigarrow & \text{final-income choice,} \\
\text{budget constraint} & \rightsquigarrow & \text{translated simplex } v + \Delta_B, \\
\text{simplex geometry} & \rightsquigarrow & \text{least-majorized point,} \\
\text{Schur objectives} & \rightsquigarrow & \text{objective collapse at water-filling.}
\end{array}$$

Geometric reduction. The planner starts with transfers. First, for the named objectives, unused budget can be spent without hurting the objective, so we may restrict to exact-budget transfers $x \in \Delta_B$. Second, transfers matter only through final incomes, so we pass from x to $y = v + x$. The feasible final-income set is therefore the translated simplex $v + \Delta_B$. On this simplex, water-filling is the least-majorized point, and Schur-respecting objectives all select it.

Two moves. Two simple moves turn the planner's problem into geometry. The first is a reduction on transfers: for the named monotone objectives, some optimum exhausts the budget, so we may restrict from weak-budget transfers to the exact-budget simplex Δ_B. The second is a change of viewpoint: the objectives evaluate final incomes, not transfers themselves, so an exact-budget transfer $x \in \Delta_B$ is viewed through its image $y = v + x$. Combining the two moves gives the translated simplex $v + \Delta_B$. The rest of the note studies this simplex.

$$\begin{array}{ccccccc}
\begin{array}{c}\text{planner with budget } B \\ \text{objective dilemma}\end{array} & \Longrightarrow & \begin{array}{c}\text{choose allocation } x \\ \text{many feasible transfers}\end{array} & \Longrightarrow & \begin{array}{c}\text{choose final income } y = v + x \\ \text{exact-budget boundary}\end{array} & \Longrightarrow & v + \Delta_B \\
 & & & & \Uparrow & & \\
\text{MaxMin} \quad \text{Gini} \quad \mathrm{FGT}_{\alpha,r} & \Longrightarrow & \text{Schur-respecting class} & \Longrightarrow & \text{least-majorized point} & & y^* = v + x^*
\end{array}$$

$$\boxed{\text{objective choice collapses: water-filling is forced by the geometry of } v + \Delta_B}$$

Figure 2: Storyline of the note. The planner begins with an objective-choice problem over transfers. Passing to final incomes and exact-budget transfers turns the feasible set into the translated simplex $v + \Delta_B$. On that simplex, water-filling is the least-majorized point. Since the objectives of interest are Schur-respecting, they all select the same final-income vector.

D A Greedy View of Objective Collapse

Algorithmic storyline. There is another way to see why the objectives agree. The key point is local, not just global: before proving that the objectives share the same final optimizer, we can see that they make the same next-dollar decision at every intermediate vector. Imagine allocating the budget one marginal dollar at a time, or one infinitesimal unit in the continuous model. At an intermediate final-income vector $\boldsymbol{y}$, ask where the next dollar should go. The comparison keeps the total transfer fixed; it only changes which coordinate receives the same marginal amount. If one agent is poorer than another, giving the dollar to the poorer agent closes an income gap, while giving the same dollar to the richer agent leaves that gap open, or makes it larger. Thus every objective considered here ranks the poorer-target move above the richer-target move. MaxMin, Gini minimization, FGT minimization, and the Schur-respecting objectives all choose the same target set,

$$i \in \arg\min_j y_j.$$

The objectives may disagree about how to score the whole allocation, but they agree on this next-dollar comparison. They remain cardinally different: they assign different numerical scores to allocations and encode different normative criteria. The greedy argument uses a weaker fact. At each intermediate vector, the planner only compares local moves: should the same next dollar go to a poorer income class or to a richer one? The objectives may measure the gain differently, but they rank the direction the same way. The collapse is local and ordinal: different cardinal objectives induce the same ordering over the relevant marginal choices.

The same marginal decision can be compressed into larger steps. Once the objectives agree that the next dollar goes to the poorest income class, we can keep making that same choice until the choice would no longer be the same. This means raising the currently poorest agents by the largest nontrivial amount that does not change the income order, or by the remaining budget if it runs out first. The larger step is only a compression of repeated marginal choices. It stops exactly when the lowest income level reaches the next distinct income level. At that point, the set of poorest agents has changed, and the same argument applies again.

When several agents are tied at the minimum, the relevant target is the whole lowest-income class, not an arbitrary representative. The compressed step removes tie-breaking: instead of choosing among tied poorest agents one unit at a time, it raises the tied class uniformly until the next breakpoint. This keeps the path canonical. Raising the class together is the order-preserving version of repeating the same next-dollar decision among tied poorest agents.

Greedy reduction. This gives a black-box reduction from the objectives to a common greedy rule. Start from $\boldsymbol{y} = \boldsymbol{v}$. At each step, identify the agents with minimum current income and give them the next marginal dollar. Equivalently, raise the entire current lowest-income class until it either catches the next income class or exhausts the remaining budget. Thus the objectives do not independently search the feasible set: at every intermediate vector, they make the same next-dollar comparison and choose the same equalizing move,

send the next dollar to the current poorest income class.

Repeating this rule gives the water-filling path. Starting from $\boldsymbol{y} = \boldsymbol{v}$, the current lowest-income class is raised until it reaches the next distinct income level, or until the budget is exhausted. If a tie is created, the tied class is raised together in the next step.

Formal step size. For completeness, the compressed step can be written explicitly. At an intermediate vector $\boldsymbol{y}$, let $L(\boldsymbol{y}) = \arg\min_i y_i$ be the current lowest-income class, and let B_{rem} be the remaining budget. Let $y_{(1)}$ be the lowest income level and $y_{(2)}$ the next distinct income level, with $y_{(2)} = +\infty$ if all coordinates are tied. The gap to the next breakpoint is $\delta(\boldsymbol{y}) = y_{(2)} - y_{(1)}$. Since each dollar is assigned to the current lowest-income class until the next tie, the common step size is

$$\eta(\boldsymbol{y}) = \min\left\{\delta(\boldsymbol{y}), \frac{B_{\text{rem}}}{|L(\boldsymbol{y})|}\right\}.$$

Each coordinate in $L(\boldsymbol{y})$ is raised by $\eta(\boldsymbol{y})$, and the step spends $|L(\boldsymbol{y})|\eta(\boldsymbol{y})$. If the second term is smaller, the budget is exhausted. If $\delta(\boldsymbol{y})$ is smaller, the lowest-income class catches the next income level and the algorithm repeats from the new vector.

Termination and interpretation. The process produces the water-filled vector $\boldsymbol{y}^*$. Each nonterminal compressed step merges the current lowest-income class with the next distinct income level, so in the continuous model there are at most n such steps, or fewer if the budget is exhausted earlier. In a discrete instance with indivisible unit transfers and $B \in \mathbb{Z}_+$, the literal dollar-by-dollar implementation takes at most B unit steps; the compressed implementation groups those identical marginal choices into at most n leveling steps. The common object is the final-income vector $\boldsymbol{y}^*$; the induced transfer is $\boldsymbol{x}^* = \boldsymbol{y}^* - \boldsymbol{v}$. Thus the objectives agree for a stronger reason than shared optimality ex post: along the whole path, they assign the same marginal dollars to the same lowest-income targets and converge to the same allocation $\boldsymbol{x}^*$, equivalently the same final-income vector $\boldsymbol{y}^*$.

This is the pathwise form of objective collapse. The agreement is not an endpoint coincidence, and it is not a claim that the objectives are normatively identical. The objectives do not merely intersect at the same optimizer; they reveal the same marginal rule along the route to it. This greedy picture is the algorithmic shadow of the least-majorization theorem: each step performs the locally most equalizing order-preserving move, and the endpoint is the unique least-majorized feasible vector. In this sense, the direct-transfer simplex converts normative disagreement over objectives into agreement over the same equalizing marginal policy.