The Welfare Representation of Envy-Free Choice*

Niels Boissonnet
Department of Economics, Center for Mathematical Economics (IMW), Joint Institute for Individualisation in a Changing Environment (JICE), Bielefeld University

Yves Breitmoser
Department of Economics, Center for Mathematical Economics (IMW), Joint Institute for Individualisation in a Changing Environment (JICE), Bielefeld University

October 2026 · Working paper

How can a planner compare people's welfare so that the comparison respects envy freedom, the requirement that neither person prefer the other's allocation? Conventional equivalents preserve individual preferences, but their reference assignments may give the wrong fairness judgment. We characterize fairness-compatible equivalence-based normalizations for ordered two-person efficient allocation paths with adequate reference coverage. The appropriate reference for assessing another person's package is that recipient's efficient assignment path; the resulting Contract Curve Equivalents reproduce both directions of envy. Different envy cutoffs generally require two directed pairs of indices, while relative-cost variation restricts their common scaling and, under further invariance conditions, identifies the physical references used. We also show why fixed own-welfare reports cannot generally support consistent efficient no-envy choice across problems. A compensated benefit reform may preserve all attainable own-utility profiles while changing or eliminating efficient envy-free allocations. Thus policy evaluation under no-envy requires information about how each person values the other's package.

Key words: fair social choice; no-envy; interpersonal welfare comparisons; social welfare functions; equivalence-based well-being measures; risk sharing
JEL codes: D60; D63; D71; D81

1. INTRODUCTION

Economists often assign a numerical welfare level to each person before comparing policies. Its reference matters when the comparison must respect a fairness principle. Consider an efficient transfer that improves one recipient's package at another's expense. Under *envy freedom*, each person must prefer her own package to the other's actual package. The planner therefore needs more than reports of what each receives from her own assignment. Such comparisons matter when benefits or risks differ in composition, as research on in-kind transfers and public-housing assignments illustrates [Currie and Gahvari, 2008, Waldinger, 2021]; those institutions need not themselves adopt no-envy. Fair income taxation and generalized welfare weights address related decisions under different evaluative criteria [Fleurbaey and Maniquet, 2006, Saez and Stantcheva, 2016]. We characterize the equivalence-based normalizations that represent no-envy on ordered two-person efficient paths, identify when two directed index pairs are necessary, and show why changes in package composition can matter even with unchanged own-welfare possibilities.

* We thank Gerrit Bauch, Juergen Eichberger, Sebastian Fehrler, Jurek Preker, Frank Riedel, and Karl Schlag for helpful comments. We also thank audiences at the Economic Theory Lunch Seminar at Bielefeld University, the economics seminar at the University of Bremen, and the annual meeting of the VfS Committee for Economic Theory. Financial support from the Verbundinitiative InChange, funded by the Ministry of Culture and Science of the State of North Rhine-Westphalia, is gratefully acknowledged.

The problem appears even when both people agree on the fair allocation. Suppose they divide the promised amounts of two payments, each delivered with probability $1/2$. Person 1 is risk neutral and person 2 has square-root utility. Since each prefers a larger promised amount, the equal split is envy free. Now have a planner maximize the *smaller* of their reported welfare numbers, a maximin rule. If those reports are *certainty equivalents*, the sure payments each person considers as good as her lottery, person 1 receives only one third of the total promise. If instead they are *probability equivalents*, winning probabilities for a lottery with the same fixed prizes, person 1 receives about 0.618. In the first case person 1 envies; in the second person 2 envies. The physical lotteries and the maximin rule have not changed. The recommendations differ because sure-cash and probability references produce different welfare comparisons, neither of which tests whether each person prefers her own *contingent payment*. We derive the calculations in Section 2.1.

The reference in an interpersonal welfare measure is a substantive choice. Pazner and Schmeidler [1978] study egalitarian equivalence using a common reference bundle; Fleurbaey and Maniquet [2011, 2017, 2018a] characterize preference-respecting measures whose reference rays or budgets express fairness requirements, with implications for taxation [Fleurbaey and Maniquet, 2018b]. Piacquadio [2017] derives opportunity-equivalent utilitarianism from fairness axioms, while Fleurbaey and Zuber [2021] characterize a normalization using transfer and risk-aversion requirements. Their criteria do not themselves establish a representation of *both directed no-envy comparisons* between actual packages. We therefore ask which references the no-envy principle requires, rather than declaring these existing measures unfair. Related analyses concern common scales, envy intensity, multidimensional disadvantage and reference consumption [Caradonna and Raines, 2026, Bosmans and Öztürk, 2018, De Rock and Moramarco, 2026, Maniquet and Moramarco, 2026].

No-envy gives a direct way to judge whether a welfare comparison respects a specified fairness principle. At an allocation, person 1 compares her package with person 2's using her preferences; person 2 makes the reverse comparison using hers. No common numerical utility scale is needed. The criterion is central to fair-allocation research [Foley, 1967, Varian, 1974], while a recent working paper by Schlag [2026] examines its normative foundations. We take no-envy as the fairness requirement to be represented, not as a complete account of justice: rights, needs, and responsibility can matter even when neither person envies. Nor does the lottery example prove that *every* fixed welfare measure fails. To establish that claim within a specified information setting, we must say what the planner observes and how consistently she must choose.

Suppose the planner sees only two numbers reporting what the recipients obtain from *their own* assignments, and the indices producing those numbers remain fixed as opportunities change. Our first theorem constructs nearby allocation problems in which the same two welfare reports are feasible, but efficient no-envy uniquely requires choosing the first report in one problem and the second in the other. A rule based on those reports must reverse its ranking of unchanged alternatives, violating the weak axiom of revealed preference (WARP). This differs from the previously recognized difficulty of expressing reciprocal envy through a single interpersonal ordering [Fleurbaey and Maniquet, 2014, pp. 7–8], and from impossibilities for fair collective choice on broader domains [Sakamoto, 2013, Fleurbaey and Tadenuma, 2014]. The specific obstacle here is missing information about how each person values the *other's* assignment. A planner who observes the packages and both recipients' preferences can test envy directly; the question is how to preserve those comparisons in welfare measures.

A continuous transfer or resource share gives a familiar economic foundation for the positive construction. When increasing the coordinate improves person 1's package and worsens person 2's, and *all* efficient allocations lie on that path, the decision has an *ordered contract curve*, provided both evaluators rank each recipient's packages in the same direction. Neither an adjustable tax rate nor an insurance premium automatically has this property: incentives, resources, or the composition of packages may change nonmonotonically. We also need adequate reference coverage so that each actual package has an indifferent equivalent on each recipient's path. The separable multi-good resource economy proved below satisfies both requirements, as does the running cash–quality economy. These primitive examples establish a meaningful domain of continuous distributive problems without claiming that every one-dimensional policy belongs to it [Moulin, 2017].

The first positive discovery identifies the *physical reference* for each missing welfare comparison. To test person 1's envy, find the package on person 2's efficient assignment path that person 1 values as highly as her own, then compare its position with person 2's actual position. Person 2's test uses person 1's path instead. These indifferent positions are *Contract Curve Equivalents* (CCEs). Each directed pair makes both evaluators assign the same reported level to any package on the recipient's reference path. The construction consequently represents the corresponding no-envy test, and the common-level requirement characterizes its numerical scale up to a shared increasing relabeling. This is a characterization of suitable normalizations, not the claim that one arbitrary reference makes all social comparisons fair.

The second positive discovery explains why one such pair generally cannot do both jobs. Person 1's no-envy test sets a *lower* cutoff on the efficient policy coordinate; person 2's sets an *upper* cutoff. Along the curve, person 1's own welfare rises and person 2's falls. A single pair of preference-representing indices can therefore cross only once, so it reproduces both directed envy boundaries on the efficient curve exactly when the cutoffs coincide. When the cutoffs differ, two directed pairs are generally required. The running cash–quality example has a nondegenerate fair interval and illustrates that case. This one-pair exception concerns comparisons on the efficient curve, not agreement on every inefficient allocation.

Agreement with one observed envy boundary would make this construction possible, but not yet necessary. To distinguish a genuine characterization from an accidental fit, we vary relative production costs while holding the *recipient's reference path* and both preferences fixed. Different packages then become the binding envy comparison. An index pair that itself remains fixed across these cost experiments must agree on the entire tested path and is consequently a shared increasing transformation of the CCE pair. Further common-reporting, dense-testing, and equality-invariance conditions identify the *physical reference assignments actually used* when candidate paths and numerical scales may vary. They do not identify untested extensions. The distinction between representing preferences and fixing an interpersonal numerical scale also arises in Morreau and Weymark [2016].

Why construct welfare indices when the planner could test no-envy directly? Direct comparisons determine the fair *set*, but selecting a policy within it requires another social judgment. In the cash–quality example, the fair shares range from 0.4189 to 0.4384; choosing 0.43 requires additional preferences. We characterize interval choice and show when calibrated CCE mixtures and maximin represent the observed selection. Disconnected menus and cross-economy choices impose further consistency conditions. These are restrictions on social choice, not consequences of no-envy. Sher [2024] examines policy-ranking inconsistency in a different framework, which our results do not claim to resolve.

The normalization matters when policies change physical packages. If person 2 receives better service, person 1 may need a higher resource share before she stops envying; person 2 may also tolerate a larger share for person 1 before she envies. Thus both cutoffs can move rightward even with unchanged preferences and fairness requirements. A welfare index tied to the previous recipient path may then give the wrong fairness comparison. This comparative statics concerns the packages being evaluated, not a change in anyone's moral beliefs.

A separate linear cash–service economy isolates the missing information. Let $U_1(m, q) = m + 2q$ and $U_2(m, q) = m + q$. Initially each receives $(0.5, 0.5)$; a compulsory reform gives person 1 $(0.3, 0.6)$ and person 2 $(0.4, 0.6)$. Both recipients retain own utilities 1.5 and 1, and public spending remains 2.5 at service price 1.5. Yet person 1 values person 2's new package at 1.6, above her own utility 1.5, and begins to envy. Unlike the nonlinear running example, this model can preserve the *entire* feasible own-utility possibility set while changing or eliminating the efficient envy-free set. Retaining the original packages as options changes that conclusion. The implication is conditional on no-envy, not a judgment by every fairness criterion.

The analysis also gives a conditional policy-evaluation procedure: elicit or calculate each person's value of the other's actual package, locate the two envy cutoffs, and apply a separate social ranking if several fair policies remain. After packages change, own-welfare reports alone do not establish that the cutoffs remain fixed. In the linear cash–service model, known preferences, budget, service price and total service, together with own-utility reports, recover the two missing envy gaps at a budget-exhausting allocation. In settings without those primitives, the analyst may need direct comparisons of recipients' packages. This is not a formula for an optimal tax or premium.

Throughout, preferences and the no-envy criterion are fixed. Actual judgments can also respond to

endowments and beliefs [Charité et al., 2022, Alesina and Angeletos, 2005]. Those findings motivate studying reforms but do not establish that recipients use CCEs or that our model predicts changing moral judgments. We characterize how the *same* envy criterion evaluates changed packages.

Our contribution is a characterization of no-envy-compatible welfare normalization, not a new fairness axiom or a universal policy ranking. Ordered contract curves and Range delimit the domain, rich cost tests discipline the references, and additional social-choice assumptions govern selection. The benefit and risk-sharing results establish why the distinction matters economically.

2. FROM SIMPLE NO-ENVY TO THE INFORMATION PROBLEM

2.1. One-dimensional problems and the contingent-payment example

Before introducing welfare numbers, we can ask directly whether a person envies. An *assignment* $x \in X$ is the package one person receives; an *allocation* $g = (g_1, g_2) \in X^2$ assigns a package to each person. Person 1 must weakly prefer g_1 to g_2, and person 2 must weakly prefer g_2 to g_1:

$$g_1 \succsim_1 g_2 \qquad \text{and} \qquad g_2 \succsim_2 g_1.$$

The first inequality is person 1's comparison using *her* preferences. The second is person 2's comparison using *hers*. Neither test asks whether the two people have equal utility numbers, and neither requires the planner to choose a common cardinal scale.

With one desirable good, the conclusion is immediate: whoever receives less prefers the other's larger quantity, so only an equal split is envy free. If both receive payments delivered with the *same* probability, they likewise prefer the larger promised payment. Now suppose the planner replaces each risky payment by an equally preferred but different reference, such as sure cash. The actual promises are unchanged, but the reported welfare levels can lead a familiar social rule to a different split.

Let a payment of m arrive with probability $p \in (0, 1)$, and let non-delivery yield zero. Denote this lottery by $H_p(m)$ and a sure payment of m by δ_m:

$$H_p(m) = p\delta_m + (1 - p)\delta_0, \qquad S(m) = \delta_m.$$

The people divide a total promised payment $M > 0$, with the same delivery event for both. If person 1 is promised m, person 2 is promised $M - m$, so the feasible allocations are:

$$G_p = \left\{ \left(H_p(m), H_p(M - m) \right) : m \in [0, M] \right\}.$$

For a fixed p, a strictly increasing expected-utility evaluator prefers a higher promised payment. A recipient promised less than half therefore envies the other, whatever their risk attitudes. The only envy-free allocation gives both the promise $M/2$.

Keep this lottery problem and the maximin rule fixed, but report each person's *certainty equivalent*, the sure payment she values equally to her lottery. Let $p = 1/2$. Person 1 is risk neutral, with $v_1(x) = x$, and person 2 has square-root utility $v_2(x) = \sqrt{x}$. For $H_{1/2}(m)$, the sure equivalents are $m/2$ and $m/4$, respectively: person 2 demands more compensation for bearing risk. If person 1 receives fraction t of the total promise, maximin equates the two reported levels, $tM/2 = (1 - t)M/4$, which gives $t = 1/3$. Both reported sure-cash values are equal at this split, but person 1 has the smaller *promised lottery* and therefore envies. Nothing is wrong with either certainty equivalent as a representation of individual preference. Equality of those equivalents simply does not compare the actual contingent payments on a common basis.

Try another familiar measurement. A *probability equivalent* asks for the winning probability with which a person would accept the common prizes 0 and M instead of her actual lottery. Person 1 finds $H_{1/2}(tM)$ equivalent to winning M with probability $t/2$; person 2 assigns the lottery $H_{1/2}((1 - t)M)$ an equivalent probability $\frac{1}{2}\sqrt{1 - t}$. Maximin of these probability reports now requires $t = \sqrt{1 - t}$. This time $t = (\sqrt{5} - 1)/2 \approx 0.618$, so person 2 receives the smaller promised payment and envies. Sure-cash equivalents recommend one third of the promise for person 1, and probability equivalents recommend about 62 percent, although the fair split of the actual lotteries remains one half. Both indices correctly rank each recipient's lotteries. Their failure as no-envy tests comes from measuring them against *different reference assignments*

rather than comparing the contingent payments being exchanged. The calculation does not use the later Range Assumption.

With preferences, feasible lotteries, and maximin held fixed, the two reference choices nevertheless recommend different envious allocations. Changing the aggregator alone does not explain which comparison is appropriate. Before aggregating, the planner needs measures that express both people's valuations of the *same recipient's actual assignment* on a common reference.

2.2. Why fixed welfare indices of own assignments are insufficient

The lottery example compared two familiar measures. Could *some other fixed pair* of preference-representing indices implement efficient no-envy whenever economic opportunities change? Let person i's complete, transitive preference $\succsim_i$ be represented by a continuous utility U_i. An allocation problem has a compact feasible set $G \subseteq X^2$. We write $\text{PE}(G)$ for its Pareto-efficient allocations and $\text{PEEF}(G)$ for those efficient allocations that also pass both no-envy tests. Checking a single allocation g requires four evaluations:

$$U_1(g_1), \quad U_1(g_2), \quad U_2(g_1), \quad U_2(g_2).$$

The first and last numbers measure what the two people receive. The middle two answer the questions welfare reports may omit: how does person 1 value person 2's package, and how does person 2 value person 1's?

Now suppose the planner receives only one number for *each person's own package*. Let e_i represent person i's preferences, with both e_i's fixed as opportunities change. At allocation g she observes the pair $W_e(g)$; for an entire problem she observes the set $W_e(G)$ of such pairs:

$$W_e(g) = \big(e_1(g_1), e_2(g_2)\big), \qquad W_e(G) = \{W_e(g) : g \in G\}.$$

A report such as $(2, 1.5)$ says nothing about whether person 1 prefers the physical package behind person 2's report. Even knowing every attainable report in $W_e(G)$ need not reveal those packages. This is a precise limitation on what the planner observes, not a claim that all welfare-based choice rules must discard physical assignments.

Suppose two reports ξ and ξ' are feasible in each of two economies, but only an allocation reporting ξ is envy free in the first, and only one reporting ξ' is envy free in the second. Fair choice would then require the planner to select ξ over ξ' in one economy and reverse that selection in the other, even though both remain available. A single-valued menu-independent rule cannot do this and satisfy the weak axiom of revealed preference (WARP): if a report is selected when another is available, the rule cannot select the latter instead from another menu where both are available. We first exhibit the reversal with actual packages; the theorem then shows that it can occur arbitrarily near equal allocations for *any* fixed preference-representing indices.

Take cash and quality packages (m, q), with the following preferences:

$$U_1(m, q) = m + q, \qquad U_2(m, q) = m + \frac{q}{2}.$$

The equal package $x = (1, 1)$ is worth 2 to person 1 and 1.5 to person 2. A different package $y = (0.5, 1.7)$ is worth 2.2 to person 1 but only 1.35 to person 2: person 1 prefers y, while person 2 prefers x. Using this disagreement, form two line-segment feasible sets. Each contains one equal-package allocation, its unique envy-free choice, and a second allocation that reports the same own-utility vector as the *other* problem's equal-package allocation:

Problem	Allocation (g_1, g_2)	No-envy?	Utility vector (U_1, U_2)
G, center	$((1, 1), (1, 1))$	yes, uniquely	$\xi = (2, 1.5)$
G, end	$((1.1, 1.1), (0.9, 0.9))$	no, person 2 envies	$\xi' = (2.2, 1.35)$
H, center	$((0.5, 1.7), (0.5, 1.7))$	yes, uniquely	$\xi' = (2.2, 1.35)$
H, end	$((0.4, 1.6), (0.6, 1.8))$	no, person 1 envies	$\xi = (2, 1.5)$

In the first problem G, only the equal allocation at x is envy free, so a fair rule selects $\xi = (2, 1.5)$ rather than the available $\xi' = (2.2, 1.35)$. In H, only the equal allocation at y is envy free, so the rule must select ξ' despite the continued availability of ξ. The reported alternatives are unchanged, but their underlying physical assignments are not. A rule using only the reports therefore reverses its revealed preference. The theorem shows why no alternative fixed preference-representing scales eliminate the problem.

Theorem 1 (Fixed own-welfare reports cannot generally implement no-envy)**.** *Let x be interior to $X \subseteq \mathbb{R}^d$, with both preferences continuous and strictly monotone nearby; every neighborhood contains an assignment preferred to x by one person and dispreferred by the other.*

Arbitrarily near (x, x), compact Pareto-efficient line-segment problems G, H have oppositely ordered recipient paths. For any fixed preference-representing pair e, $W_e(G)$ and $W_e(H)$ share distinct reports ξ, ξ', but their unique efficient envy-free outcomes yield ξ and ξ', respectively. Thus no single-valued, menu-independent rule based only on these reports can implement efficient no-envy throughout a neighborhood while satisfying WARP.

The reversal comes from disagreement over the change from x to y. Person 1 favors it and person 2 opposes it. Near these packages we can adjust the other feasible assignments so that two own-utility reports recur, while each economy has a different unique equal-package, envy-free allocation. When both preferences are differentiable and their gradients are linearly independent, we can make these adjustments while keeping total resources constant along each line segment. The resource-balanced construction is proved in the Additional Material. With one desirable good, by contrast, both recipients prefer a larger quantity; no such opposing local rankings are available.

It is well known that own-utility information can miss distributive judgments [Sen, 1979, Roemer, 1996]. Here the specific statement is stronger than the observation alone: *arbitrarily close* ordered problems produce opposite uniquely fair choices between the *same* own-welfare reports, for every fixed utility representation. WARP detects the conflict. Moulin and Thomson [1988] consider Pareto efficiency, resource monotonicity, and no-domination, and Sakamoto [2013] study collective rationality and envy-free selection on richer domains. Our theorem isolates the information withheld from this planner; it does not prohibit a planner with direct access to preferences and assignments from checking envy.

The impossibility requires indices that remain fixed across problems and a planner who sees only their reports. With the assignments and both people's rankings in hand, she could test envy directly. Our positive construction expresses those cross-package comparisons numerically by using the efficient assignments of the *recipient being compared* as references.

2.3. Ordered contract curves

Many distributive decisions have a continuously adjustable transfer, budget share, or cost contribution. When both people prefer more cash, dividing a fixed cash budget with no other instruments gives this ordering: increasing person 1's share improves her package and reduces person 2's, and the feasible splits exhaust the efficient allocations. The same reasoning can apply to an efficient policy path with several goods or risky contracts, provided both people agree about the direction in which *each recipient's* package improves. It does not follow merely because the government chooses a single tax rate or premium: changing that rate may change labor supply, aggregate resources, and both people's packages in ways that violate the required ordering. We therefore assume more than one-dimensionality. An *ordered contract curve* is the full Pareto-efficient set, parametrized so that both recipients rank each person's assignment path in the same direction and the two envy comparisons each have a cutoff. The definition makes these conditions precise.

Definition 1 (Ordered contract curve)**.** A feasible set G has an ordered contract curve if there are continuous paths

$$C_{G,1}, C_{G,2} : [0,1] \to X$$

that both people strictly rank in the increasing direction and for which

$$\gamma_G(t) = \big(C_{G,1}(t), C_{G,2}(1-t)\big), \qquad t \in [0,1],$$

is a continuous bijection from $[0,1]$ onto $\mathrm{PE}(G)$. At the endpoints, require

$$C_{G,2}(1) \succsim_i C_{G,1}(0), \qquad C_{G,1}(1) \succsim_i C_{G,2}(0), \qquad i = 1, 2.$$

The two paths in the definition describe *individual packages*: $C_{G,1}(t)$ is a package person 1 could receive efficiently, while $C_{G,2}(r)$ is one person 2 could receive. The curve $\gamma_G(t)$ specifies both packages at once, so raising t improves person 1's package in both people's eyes while worsening person 2's. The endpoint

comparisons ensure that each person's ranking of the two packages switches somewhere along the curve. Continuity and strict ordering then give one cutoff for each envy test, although these cutoffs need not coincide. The parameter t is a contract-curve *coordinate*; only in specific resource models is it a physical share.

The ordering also arises with several divisible resources. Under strictly increasing, strictly concave, additively separable utilities, raising person 1's Pareto weight increases her efficient resource assignment and reduces person 2's. Both people therefore rank each recipient's packages in the same direction. The proposition shows that the efficient paths extend continuously from receiving nothing to receiving the whole endowment. Because the two paths cover those extremes, any feasible package has a unique indifferent position on either path, which Section 3.1 will use to define CCEs. Thus the model gives a *sufficient economic foundation* for both ordered choice and reference-path coverage, rather than merely assuming that a continuously adjustable instrument has those properties. The Online Appendix proves the claim and gives an efficient envy-free existence argument [Varian, 1974].

Proposition 1 (Separable resource allocation generates an ordered contract curve)**.** *Let $B \in \mathbb{R}^S_{++}$ and $x_1 + x_2 \leq B$, with utilities*

$$U_i(x_i) = \sum_{s=1}^{S} a_{is} u_{is}(x_{is}),$$

where $a_{is} > 0$ and each $u_{is} : \mathbb{R}_+ \to \mathbb{R}$ is continuous, strictly increasing, and strictly concave. The efficient set is an ordered contract curve; both paths run continuously from 0 to B, and each feasible assignment has a unique indifferent equivalent on either path for either evaluator.

We can locate the two cutoffs without choosing welfare indices. At coordinate a_G, person 1 is indifferent between her own package and person 2's; at b_G, person 2 is indifferent. Thus they solve two different *physical-package comparisons*:

$$C_{G,1}(a_G) \sim_1 C_{G,2}(1 - a_G), \qquad C_{G,2}(1 - b_G) \sim_2 C_{G,1}(b_G).$$

Below a_G, person 1 prefers person 2's package and envies. Above b_G, person 2 prefers person 1's package and envies. Both tests therefore pass exactly when $t \geq a_G$ and $t \leq b_G$. The efficient envy-free set is

$$\mathrm{PEEF}(G) = \begin{cases} \gamma_G([a_G, b_G]), & a_G \leq b_G, \\ \varnothing, & a_G > b_G. \end{cases}$$

If $a_G \leq b_G$, every efficient allocation between the cutoffs is envy free; if $a_G > b_G$, none is. Notice that the two cutoff equations use only the recipients' preferences for actual packages. Welfare indices are a way to *report* these judgments, and a social choice rule is needed only if several fair allocations remain.

For one divisible good the two cutoffs coincide at equal division. With cash and service quality they can differ. Suppose one unit of input can be split between two production processes. Input s for person 1 produces $C_1(s) = (s, s)$, while input r for person 2 produces $C_2^0(r) = (r, r^2)$, with $s + r \leq 1$. Both people value more cash and more service, but person 1 values service more strongly:

$$U_1(m, q) = m + 2q, \qquad U_2(m, q) = m + q.$$

Using all input is efficient because both people prefer more of their own cash and service. If person 1 receives fraction t of the input, person 2 gets $1 - t$, so the efficient allocation is $\gamma^0(t) = \big(C_1(t), C_2^0(1 - t)\big)$.

Each person must evaluate *both* packages using her own preferences. Person 1 values her own at $3t$ and person 2's at $(1 - t) + 2(1 - t)^2$. Person 2 values her own at $(1 - t) + (1 - t)^2$ and person 1's at $2t$. Equating the two values separately for each evaluator gives the two envy cutoffs:

$$a_0 = 2 - \frac{\sqrt{10}}{2} \approx 0.4189, \qquad b_0 = \frac{5 - \sqrt{17}}{2} \approx 0.4384.$$

Hence any input share for person 1 between about 0.4189 and 0.4384 is efficiently envy free. Because these are different shares, no-envy alone does not select one allocation. The example also supplies the coverage needed for CCE measurement: both recipient paths run from the common worst package $(0, 0)$ to the common best package $(1, 1)$, and both evaluators strictly rank intermediate packages. Every package in $X = [0, 1]^2$ therefore has exactly one equally preferred position on either path. We next use those positions to report the two comparisons numerically.

3. REPRESENTING THE TWO DIRECTIONS OF ENVY

The fixed own-welfare reports omit the evaluations that each person makes of the other's package. We now construct welfare measures that report those comparisons by referring to the *efficient packages of the recipient being compared*. The purpose is not to choose an attractive normalization, but to characterize one that preserves the two no-envy tests. Each direction has an index pair of its own; the representation theorem determines when the same pair can serve both.

3.1. The contract-curve-equivalent construction

Suppose person 1 wants to compare her package with person 2's. Converting her own package to sure cash would not by itself tell her whether she prefers person 2's *actual* package. Instead, move along person 2's efficient assignment path until she finds a package she regards as equally good as her own. She can compare the resulting position directly with person 2's actual position on that same path. If person 2 occupies the higher position, person 1 envies. For person 2's envy we repeat the calculation on person 1's path. We first need to make sure these indifferent packages exist.

Fix one allocation problem H and write its ordered efficient curve as $\gamma(t) = \big(C_1(t), C_2(1-t)\big)$ for $t \in [0,1]$, suppressing the problem index. The path C_1 lists the packages person 1 can receive efficiently, from worst to best; C_2 does the same for person 2. Both people agree that moving upward on either path gives its recipient a better package. Person 1 must measure *her own* package on C_2 to compare it with person 2's, while person 2 measures hers on C_1. The superscript on the coming equivalent will record which recipient supplies this comparison path.

A reference path cannot measure a package that one evaluator likes more than every package on the path or less than every package on it. We therefore specify a domain D containing the two complete efficient recipient paths and require that *each person* find a unique equally preferred package on *each path* for every package in D. The Range Assumption states exactly this coverage requirement.

Assumption 1 (Range)**.** *There is a domain $D \subseteq X$ containing both efficient assignment paths, $C_1([0,1]) \cup C_2([0,1]) \subseteq D$, such that, for every person $i \in \{1,2\}$, every $k \in \{1,2\}$, and every $x \in D$, there is a unique $r \in [0,1]$ satisfying*

$$C_k(r) \sim_i x.$$

Range guarantees an indifferent reference position for every package we need to compare. Because D contains *both complete paths*, it also requires each evaluator to regard their two worst endpoints as equally good and their two best endpoints as equally good. This is stronger than finding an equivalent for one observed package. It holds in the cash–quality example on $X = [0,1]^2$ and in the separable resource economy, where both paths reach 0 and the full endowment B. Range supplies reference positions, not a cardinal interpretation of differences between those positions.

We call the unique indifferent position the *contract-curve equivalent* (CCE). Its notation records both whose preferences evaluate a package and whose efficient assignments supply the reference.

Definition 2 (Contract Curve Equivalent)**.** Under Assumption 1, person i's contract-curve equivalent of assignment $x \in D$ on recipient k's path is the unique number $\pi_i^k(x) \in [0,1]$ satisfying

$$C_k\Big(\pi_i^k(x)\Big) \sim_i x.$$

Read $\pi_i^k(x)$ as "person i's equivalent of package x on recipient k's path." The number is a *position*, not an amount of utility gained. If $x = C_k(r)$ is already on that path, both evaluators report exactly r: $\pi_1^k(C_k(r)) = \pi_2^k(C_k(r)) = r$. They can therefore compare an actual recipient package using a common number without agreeing on how much utility a move of, say, 0.1 provides.

At coordinate t, person 2 receives package $C_2(1-t)$. Person 1 values her own package $C_1(t)$ as highly as a package at position $\pi_1^2(C_1(t))$ on person 2's path. Since she likes higher positions on that path, she does not envy exactly when *her indifferent position* is at least as high as *person 2's actual position* $1-t$:

$$C_1(t) \succsim_1 C_2(1-t) \quad \Longleftrightarrow \quad \pi_1^2(C_1(t)) \geq 1-t. \tag{1}$$

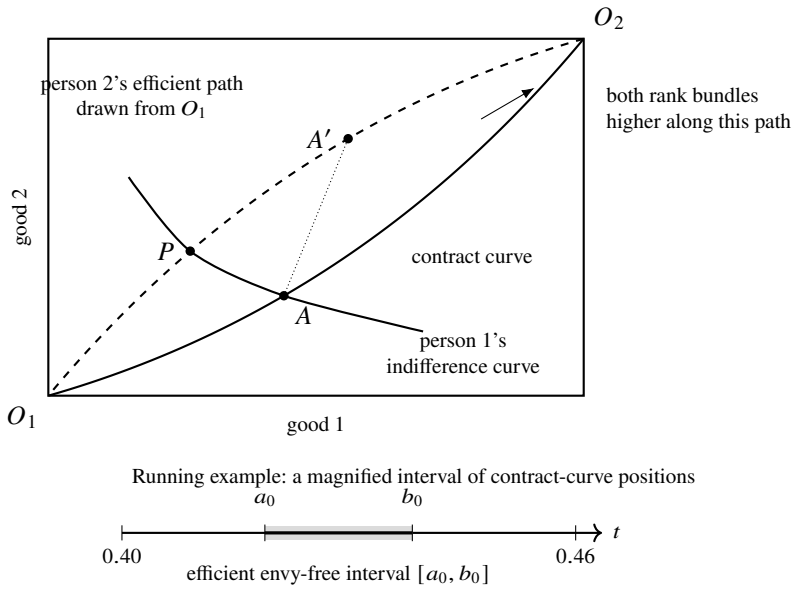


Figure 1. The two directed no-envy comparisons. In the schematic Edgeworth box, A is person 1's bundle and A' is person 2's bundle drawn from O_1. Person 1 is indifferent between A and P on person 2's efficient assignment path, but prefers A' to P, so she envies person 2. The lower panel shows the distinct equality boundaries a_0, b_0 of the cash–quality example and its efficient envy-free interval.

For person 2 we reverse the roles. She measures her own package on person 1's path, because her question is whether she would rather receive person 1's actual package:

$$C_2(1-t) \gtrsim_2 C_1(t) \quad \Longleftrightarrow \quad \pi_2^1(C_2(1-t)) \geq t. \tag{2}$$

Both inequalities now report actual preference comparisons as inequalities between numerical positions. The first uses person 2's path and binds at a_H; the second uses person 1's path and binds at b_H. They are not interchangeable: measuring person 1's envy on person 1's own path would compare the wrong reference assignments.

Return to the cash–quality economy. Person 1 must find the indifferent position on person 2's *nonlinear* service path, whereas person 2 uses person 1's *linear* path. The resulting two tests are

$$\pi_1^2(C_1(t)) = \frac{\sqrt{1+24t}-1}{4} \geq 1-t = \pi_2^2(C_2^0(1-t)), \tag{3}$$

$$\pi_2^1(C_2^0(1-t)) = \frac{(1-t)+(1-t)^2}{2} \geq t = \pi_1^1(C_1(t)). \tag{4}$$

The first inequality binds at $a_0 \approx 0.4189$, the second at $b_0 \approx 0.4384$. Consider the intermediate share $t = 0.43$. On person 2's path, person 1 finds her own package equivalent to position 0.5911, higher than the position 0.57 person 2 actually occupies, so person 1 does not envy. On person 1's path, person 2 finds her own package equivalent to position 0.44745, higher than person 1's actual position 0.43, so person 2 does not envy either.

The recipient-based references now give a numerical test of fairness: $t = 0.43$ passes both comparisons, although the two envy equalities occur at different positions. Choosing 0.43 rather than another coordinate in the fair interval remains a separate social judgment. Nor can a single recipient path replace the two references needed for the directed tests.

Figure 1 makes the distinction between a *reference* package and the *actual* other package visible. In the upper panel, P is indifferent to person 1's own assignment A, but the other recipient actually holds A'. Person 1's envy depends on comparing A' with P, not on the numerical position of A alone. The lower panel shows the economic consequence for the cash–quality example: the two directed comparisons bind at different coordinates, leaving an interval of efficient envy-free allocations.

3.2. The welfare representation

We have constructed two CCE pairs, one for each direction of envy. Could a planner use a single preference-representing pair instead? In the cash–quality example, person 1 stops envying at $a_0 \approx 0.4189$, while person 2

starts envying at $b_0 \approx 0.4384$. Along the curve, person 1's own welfare rises and person 2's falls, so one pair of preference-representing indices can cross only once. The theorem identifies the circumstances in which that single crossing suffices.

For one directed comparison, we ask for more than an index pair that happens to cross at one observed allocation. Both evaluators should give *every package recipient k could receive* the same welfare level:

$$e_1(C_k(r)) = e_2(C_k(r)) \qquad (r \in [0,1]).$$

When both indices assign the same number to $C_k(r)$, it no longer matters which evaluator reports the level of that *same physical package*. A person's report for another package x is then pinned down by the package on C_k she finds equally good: she must report that reference package's level. Range guarantees the indifferent match. The only freedom left is to relabel the common levels by one increasing function shared by the evaluators; agreement at one boundary alone would not give this conclusion.

Theorem 2 (Recipient-based welfare indices generally require two pairs)**.** *For an ordered contract curve satisfying Assumption 1 on D:*

(a) *The recipient-based CCE pairs represent the envy tests in (1) and (2).*

(b) *A continuous preference-representing pair e agrees along C_k if and only if $e_i = h \circ \pi_i^k$ on D for one continuous strictly increasing h.*

(c) *A single pair represents both preferences and envy tests on the efficient curve if and only if $a_H = b_H$.*

Part (b) makes the common-level requirement precise. Suppose both people report the same number for each package on recipient k's complete path. For any other $x \in D$, person i must report the level of the package $C_k(\pi_i^k(x))$ she finds indifferent to x. The pair is therefore a common increasing transformation of the CCEs. These shared positions permit directed preference comparisons but do not represent cardinal gains or interpersonal differences in utility [d'Aspremont and Gevers, 1977, 2002]. One accidental equality cannot fix the entire scale; Section 4.2 uses cost variation to test a succession of such equalities.

Part (c) explains why the two pairs cannot generally be merged. If $a_H < b_H$, both recipients prefer their own packages throughout the interval between the cutoffs. One pair would have to report equality there to express *both* directed tests, even while person 1's own index strictly increases and person 2's strictly decreases, which is impossible. If $a_H > b_H$, both envy throughout the interval between the cutoffs, so one pair would have to report the two numbers as strictly lower than each other simultaneously. Only when $a_H = b_H$ does one crossing suffice. Because $a_0 < b_0$, the running cash–quality example needs two pairs.

The boundary equalities resemble egalitarian equivalence when packages are consumption bundles. At a_H, person 1 is indifferent between her package and $C_2(1 - a_H)$, the package person 2 receives; at b_H, person 2 is indifferent between her package and $C_1(b_H)$. Each equality thus has a concrete common comparison bundle, but the bundles need not be the same. Pazner and Schmeidler [1978] relate efficient no-envy to egalitarian equivalence in two-person exchange; our representation keeps the two directed equalities distinct instead of presuming one common reference expresses both.

The single-pair exception applies *on the efficient curve*. Even if the two cutoffs coincide there, the same pair can fail to reproduce envy at inefficient allocations elsewhere in the feasible domain. The public-benefit economy at its critical service price supplies such an example. With more than two recipients, comparisons between several people's packages create additional restrictions, so this theorem is specifically bilateral.

Fixed own-welfare numbers can omit the evaluation of another person's package; recipient-based CCEs recover that evaluation, generally through two pairs. This raises a further question: when do familiar equivalence measures already use the required reference, and what economic variation can rule out alternatives that happen to fit one envy boundary?

4. WHEN FAMILIAR WELFARE MEASURES WORK, AND WHY RECIPIENT PATHS ARE REQUIRED

Which earlier welfare equivalents satisfy the new no-envy representation test? Common bundles, commodity rays, and reference budgets can embody meaningful fairness criteria [Pazner and Schmeidler, 1978, Fleurbaey and Maniquet, 2017]. They are not defective merely because their criteria differ from no-envy. For the comparison with recipient k, however, *both evaluators must give the same physical package received by k the*

same reported level. We first check familiar measures against that condition. We then vary relative costs so that satisfying just one observed envy equality is no longer enough, and derive the CCE form as a necessity result.

4.1. The reference-path test

Consider a risky contract received by person 2. To decide whether she envies, person 1 must evaluate *that contract*, not merely her own. If both people express its value as a sure-cash equivalent, their reports may differ because they dislike risk by different amounts. The equivalents correctly represent each person's preferences, but they do not give a common level for the same recipient package. The reference-path test detects this mismatch.

Let $R : I \to X$ be the reference packages used by some candidate measure, ordered so that both people prefer higher positions in the interval I. Person i's reference equivalent of package $x \in D$ is the position r on R she finds equally good:

$$e_i^R(x) = r \quad \Longleftrightarrow \quad R(r) \sim_i x.$$

If each indifferent reference position exists, e_i^R correctly ranks person i's packages. Yet person 1 and person 2 may give different values to *one identical risky payment* when measured in their respective sure-cash equivalents. The indices then do not agree on a package they must both assess for a directed envy test.

Pick the recipient k whose packages are being compared. Require the candidate indices to give both evaluators the same level for *every* $C_k(r)$, not just the package at one observed allocation. Under the existing coverage assumptions, the following result shows that this requirement makes the candidate pair a shared relabeling of the CCEs.

Proposition 2 (When an equivalence-based measure reproduces an envy comparison)**.** *Fix $k \in \{1, 2\}$, and let the equivalence-based indices e^R be well defined on D. They agree at every recipient-k assignment, $e_1^R(C_k(r)) = e_2^R(C_k(r))$ for all $r \in [0, 1]$, if and only if one continuous strictly increasing h satisfies*

$$e_i^R(x) = h\Big(\pi_i^k(x)\Big) \qquad (x \in D,\ i = 1, 2).$$

They then represent the directed no-envy comparison against recipient k. Agreement on both paths requires $a_H = b_H$.

The proposition requires equal reports for *the compared recipient's packages*, not agreement on an external cardinal scale. Each person's indifferent match on that path then determines her report for other packages. A pair that agrees at one observed envy boundary may still fail elsewhere; the cost experiments below supply the additional comparisons needed to test it.

Sometimes a familiar reference already describes the recipient's assignments. If people divide one good, the quantity received is a common reference. If the actual efficient contracts vary only in a shared winning probability, probability equivalents can also measure the relevant comparisons. Money and commodity-ray equivalents work when the physical assignments they reference agree in the required way. Their names alone neither qualify nor disqualify them.

For the contingent payments in Section 2.1, person 1 converts the lottery $H_{1/2}(m)$ to $m/2$ sure cash and person 2 converts *the same lottery* to $m/4$. These different levels fail the common-reference test. Fixed-prize probability equivalents also disagree for an interior promised payment. Both measures remain valid individual preference representations, but neither alone represents the actual directed envy comparison.

A more revealing example has *one familiar measure for each direction*. Suppose person 1 receives only service quality and person 2 receives only cash. At resource coordinate $s \in [0, B]$, the allocation is $g(s) = \big((0, s), (B - s, 0)\big)$. Assume each person's utility adds the value $A_i(m)$ of cash and $Q_i(q)$ of quality, so $U_i(m, q) = A_i(m) + Q_i(q)$. Suppose the two functions are continuous, strictly increasing, unbounded, and zero at zero. Given any package, we can then find an equally preferred amount of *cash alone* or of *quality alone*. These two familiar equivalents are

$$e_i^m(m, q) = A_i^{-1}\big(A_i(m) + Q_i(q)\big), \qquad e_i^q(m, q) = Q_i^{-1}\big(A_i(m) + Q_i(q)\big).$$

These reference rays continue past the finite efficient paths. Thus they can measure packages that exceed the recipient's observed efficient range; the directed comparison can be checked without claiming that the complete finite paths satisfy the global Range Assumption.

Proposition 3 (Cash and quality equivalents represent the two no-envy comparisons)**.** *Along the cash–quality frontier, cash equivalents e^m represent person 1's envy test, and quality equivalents e^q represent person 2's. If the two no-envy boundaries differ, no single continuous pair representing both individual rankings along the frontier can represent both tests.*

For person 1's envy, person 2's actual assignment is *cash only*. Both evaluators give a pure-cash package $(m, 0)$ the cash-equivalent level m, so cash equivalents report the first directed comparison correctly. For person 2's envy, the compared package is person 1's *quality-only* assignment; quality equivalents supply the shared levels. They remain distinct preference representations. For instance, if $U_2(m, q) = m + q/2$, person 2 prefers the cash endpoint $(B, 0)$ to every package on the finite quality path, so that path fails Range even though the extended quality ray supplies an equivalent.

The two boundaries need not coincide. With $U_1(m, q) = m + q$ and $U_2(m, q) = m + q/2$, person 1's equality occurs at the physical resource share $s = B/2$, and person 2's at $s = 2B/3$. All shares between these positions are envy free. Cash equivalents report one endpoint, quality equivalents the other; no single pair alone captures both comparisons.

The cash and quality examples show how familiar equivalents can represent one directed test but not necessarily both. Pazner and Schmeidler [1978] study egalitarian equivalents and maximin properties. Fleurbaey and Maniquet [1996] instead characterize compensation through reference-preference comparisons of budget sets; their references serve a different fairness requirement from the recipient paths used here. Fleurbaey and Maniquet [2014, pp. 7–8] note the conflict between reciprocal envy and a universal interpersonal well-being order, while Bosmans and Öztürk [2018, Proposition 2] use ray utility to measure directed envy intensity. We ask what each reference must be for a *qualitative* comparison with a recipient's actual package, and when one pair suffices on a particular efficient curve.

A single correct crossing does not establish a representation. Indices may report equality at the observed envy boundary yet assign inconsistent levels to other packages on the recipient's path. We next change costs so that other packages, one by one, become the binding comparison, while the recipient path used to measure them remains unchanged.

4.2. Relative-cost variation and the necessity of contract-curve equivalents

Suppose two welfare indices report equality exactly when one person stops envying at the observed allocation. This does not tell us whether they would agree at a *different* recipient package. By varying the other recipient's production cost while holding the first recipient's efficient assignment path fixed, we can make different packages on that path enter binding envy tests. A pair of indices that remains fixed and passes every test must therefore agree along the whole reference path.

In our cash–quality economy, keep the recipient paths $C_1(s) = (s, s)$ and $C_2^0(r) = (r, r^2)$. Instead of requiring $s + r \leq 1$, allow $s + \kappa r \leq 1$, where κ is the cost of producing person 2's input relative to person 1's. At $\kappa = 1$, person 2 becomes indifferent at $s = b_0 \approx 0.4384$. To make the comparison bind at a *different* package on person 1's unchanged path, take $C_1(3/8)$ and $C_2^0(1/2)$. Person 2 values both at $3/4$, since $2(3/8) = 1/2 + (1/2)^2$. At relative cost $\kappa = 5/4$, producing those two packages uses exactly the available resource. The physical indifference is now part of the efficient contract curve at $s = 3/8$, although neither person's preferences nor either recipient path has changed.

The value $5/4$ is not a normalization of welfare numbers. Person 1's package uses $3/8$ of the resource, leaving $5/8$. Producing person 2's package requires input $1/2$, so its required unit cost is $\kappa = (1 - 3/8)/(1/2) = 5/4$.

We can repeat this calculation for other packages. Pick an interior position on person 1's path, find a package on person 2's path that person 2 regards as equally good, then choose the cost that makes those two packages exhaust the resource. Reversing the recipients yields tests for the other directed comparison. The following proposition gives the domain on which any interior position can be reached.

Proposition 4 (Relative costs move the no-envy boundary through the recipient path)**.** *Fix $B > 0$ and paths $T_1, T_2 : \mathbb{R}_+ \to X$, continuous and strictly increasing for both people, with*

$$T_1(0) \sim_i T_2(0), \qquad T_1(B) \sim_i T_2(B) \qquad (i = 1, 2).$$

For $\kappa > 0$, let

$$K_{1,\kappa} = \{(T_1(s), T_2(t)) : s, t \geq 0,\ s + \kappa t \leq B\},$$
$$K_{2,\kappa} = \{(T_1(s), T_2(t)) : s, t \geq 0,\ \kappa s + t \leq B\}.$$

As κ varies, the first family preserves $T_1([0, B])$ and moves person 2's binding envy equality through every interior point of that path; the second preserves $T_2([0, B])$ and does the same for person 1.

For a chosen $s \in (0, B)$, person 2 strictly ranks each path upward, and the endpoints are equally good to her. Hence there is a unique input $q_2(s)$ such that $T_2(q_2(s)) \sim_2 T_1(s)$. Set the relative cost so that the two selected packages use exactly the resource B:

$$\kappa(s) = \frac{B - s}{q_2(s)}.$$

At this cost, both packages are feasible together without slack, and person 2 is indifferent between them. The comparison can therefore bind at *any* interior package on person 1's fixed path. If the two indices remain the same throughout these experiments, each binding comparison tests their agreement at a different package. The opposite cost family performs the analogous test against person 2's path.

The indices w_1, w_2 must be held fixed for these equalities to restrict their measurement. Normalize the unchanged paths by $C_1(r) = T_1(Br)$ and $C_2(r) = T_2(Br)$. Equality at every interior position, followed by continuity, requires the indices to agree on the entire recipient path. Preference representation then determines how they evaluate every other package $x \in D$ through its indifferent equivalent on that path. We keep D explicit: very low production costs can make packages feasible that exceed the baseline paths, and the theorem does not claim Range for all such new packages.

Theorem 3 (Relative-cost variation forces contract-curve equivalents)**.** *Under Proposition 4, let D contain both baseline paths, with a unique C_k-equivalent for each $x \in D$ and both evaluators. A continuous preference-representing pair w, fixed across $K_{k,\kappa}$, fits every binding envy equality against recipient k if and only if one continuous strictly increasing h_k satisfies*

$$w_i(x) = h_k\left(\pi_i^k(x)\right) \qquad (x \in D,\ i = 1, 2).$$

The theorem holds the welfare reports fixed while changing the cost experiment. For every chosen recipient package, the experiment supplies an allocation at which the relevant no-envy comparison binds. A successful pair must assign equal levels to the two evaluators' reports of that package. Continuity covers the endpoint packages, and each evaluator's own preference representation assigns every $x \in D$ the level of its indifferent match. Thus the pair has the CCE form, up to *one shared* increasing transformation, for this directed test. The other directed test needs the other family of cost changes.

This identifies a *numerical representation* on a given domain, not necessarily the physical packages that an economist used as her references. A candidate measure could agree on the tested portion of a recipient path and use different reference packages elsewhere. Even if the reported numbers fit every equality, knowing those numbers alone need not tell us which alternative path supplied them. The next theorem adds restrictions that identify the *used reference packages*, not merely the reported scale.

Consider an economist who changes both her reference path and the units of her reports whenever production costs change. We can no longer invoke the theorem for a *fixed index pair*. To compare her procedures across cost environments, we need one additional requirement: if both evaluators agree on the numerical level of a particular *unchanged physical package* in one experiment, they must continue to agree about that same package in the others. The common number itself may change. This is an invariance of equality judgments, not of cardinal scale.

Fix the recipient k whose packages are being compared. Imagine many cost experiments whose binding equalities involve positions arbitrarily near *every* r on that recipient's efficient path C_k. We call the tested

positions *dense*: any path position can be approached by tested ones. At r, recipient k gets $C_k(r)$ and the other evaluator j gets some $y_r \sim_j C_k(r)$. The economist may use a different candidate reference path $R_h : [0,1] \to X$ in each environment h, provided both people rank it upward. Let $q_i^h(x)$ be evaluator i's indifferent position on R_h. She reports this position through a common function

$$e_i^h(x) = \theta_h(q_i^h(x)).$$

Within one environment the same continuous increasing θ_h converts both evaluators' reference positions into reported levels. Different environments may use different θ_h's. This permits changing the units of reporting without allowing person 1 and person 2 to report the *same candidate position* in different units within the same directed test.

The equality-invariance requirement can now be stated on a concrete package. If the two evaluators agree on the report for $C_k(r)$ in one cost experiment, they must agree on that *same package* in every experiment, even if the common reported value changes. A binding test fits when y_r and $C_k(r)$ receive equal reports. Since evaluator j is indifferent between those packages, fitting the test also makes the evaluators agree on their reports for $C_k(r)$. Equality invariance then carries this information to the other environments.

Theorem A (Binding comparisons identify the used reference assignments)**.** *Under the preceding coverage, common-reporting, and equality-invariance conditions, fitting every dense boundary test is equivalent to*

$$e_i^h(x) = \varphi_h\Big(\pi_i^k(x)\Big) \qquad (x \in D,\ i = 1,2)$$

for a continuous strictly increasing φ_h in each problem, and to an increasing homeomorphism q_h onto a closed subinterval of the candidate path's parameter interval with

$$R_h(q_h(r)) \sim_i C_k(r) \qquad (0 \le r \le 1,\ i = 1,2).$$

The identification argument proceeds through four implications. A boundary test first makes evaluator j's report for y_r equal the report for $C_k(r)$. Her indifference between the two packages implies agreement of the evaluators' reports *on the same* $C_k(r)$. Equality invariance transfers that agreement to every cost experiment; dense testing and continuity extend it to each point of the recipient path. Finally, both evaluators use the *same increasing reporting function* within an experiment. Equal reports therefore correspond to a single candidate reference package that both regard as indifferent to $C_k(r)$. These packages form the identified *used subpath*; extra candidate packages not employed in any comparison remain unidentified. Without dense tests, common reporting, or equality invariance, the final inference does not follow.

The identified portion of a candidate path need not be its whole length. Every recipient package on the complete path is measured because D contains that path, but extensions beyond the actually used range are not tested. A path appropriate for *one* directed comparison also need not represent the reverse. If a pair must agree on both complete recipient paths, it must satisfy the further all-position equality

$$\pi_1^1(C_2(r)) = \pi_2^1(C_2(r)) \quad (0 \le r \le 1).$$

This all-position requirement is stronger than the condition $a_H = b_H$ for representing both directed tests *only along one efficient allocation curve*. The first restricts reported levels along an entire path; the second concerns the two envy cutoffs on a single efficient curve. Keeping them separate prevents the identification claim from exceeding its tested domain.

The two directed comparisons determine which efficient allocations are envy free, but they do not rank the allocations inside a nondegenerate fair interval. A planner's choice among those allocations requires an additional social ordering, which we next seek to represent using the same CCEs.

5. SELECTING AMONG ENVY-FREE ALLOCATIONS

Knowing the efficient envy-free set does not tell society which allocation to choose *within* it. A planner deciding a continuous transfer or benefit split may have an entire interval of acceptable policies. In the cash–quality example, every input share between $a_0 \approx 0.4189$ and $b_0 \approx 0.4384$ passes both envy tests; choosing 0.43

is an additional social judgment. Related work characterizes social orderings from ordinal preferences and fairness restrictions: Fleurbaey and Maniquet [2008] study fair social orderings, and Piacquadio [2017] studies opportunity-equivalent utilitarianism. We ask whether observed choices among fair allocations can be represented using the two directed CCE pairs. Choices from restricted intervals may reveal a single preferred share; menus with gaps can reveal further comparisons, and choices across economies must be mutually consistent. The selection requirements come from the observed social ranking, not from no-envy alone.

5.1. Revealing the peak of interval choice

Suppose society chooses share 0.43 from the whole cash–quality curve. The choice should remain 0.43 on $[0.42, 0.46]$, since it is still feasible. If instead only $[0.44, 0.46]$ is offered, a rule that prefers positions nearest 0.43 chooses 0.44. These decisions reveal a *single peak* at $\tau_0 = 0.43$, but the second interval contains no envy-free allocation because every share exceeds b_0. The peak thus describes society's preferences, whereas the two envy cutoffs describe the separate fairness requirement.

For this first choice problem, we change only *which points of one efficient curve remain feasible*. Fix a parent problem H with curve $\gamma_H([0,1])$. Every restriction $G \in \mathcal{G}_H$ retains exactly the efficient allocations $\gamma_H(J_G)$ for a nonempty compact interval J_G; any other surviving alternatives are dominated. We also include the pure interval menu $G_J = \gamma_H(J)$. We do not yet allow gaps or a restriction that turns a formerly dominated allocation into a newly efficient one. The Online Appendix gives resource constraints that produce such genuine interval restrictions.

Let $C(G) \subseteq G$ denote the allocations society selects, not the recipient paths C_k. We assume single-valued choice that depends only on the efficient set, and require that deleting unchosen alternatives cannot overturn a selection that remains feasible. This is *contraction consistency*, Sen's condition α [Thomson, 2011]. We also require the chosen coordinate to change continuously as an initial interval $[0, r]$ or a terminal interval $[r, 1]$ expands or contracts. This *one-sided continuity* rules out jumps in those two menu families; it is not a claim about every simultaneous change in both endpoints.

These requirements yield a single-peaked choice rule: from any interval, society selects the available coordinate closest to its peak [Bossert and Peters, 2009]. To see why, suppose the full curve's choice is 0.43. By contraction consistency, every interval containing 0.43 must choose it. Consider $[0, 0.42]$, entirely to its left. If it chose an interior point such as 0.41, continuity as the right endpoint approached 0.43 would require choices to pass through other points already feasible in $[0, 0.42]$. Those choices would contradict contraction consistency with the original selected 0.41. Thus any interval left of the peak selects its *right* endpoint; an interval right of the peak selects its *left* endpoint. The rule always selects the available coordinate nearest the peak:

$$C(G) = \left\{ \gamma_H\left(\text{proj}_{J_G}\, \tau_H\right) \right\}. \tag{5}$$

The projection $\text{proj}_{J_G}\, \tau_H$ equals the peak when it lies in the interval, the right endpoint when the interval lies left of it, and the left endpoint when it lies right. The fairness requirement is separate: whenever an *efficient* envy-free allocation is available, the chosen allocation must be one of them. Applied to the full curve, this puts τ_H in $[a_H, b_H]$. A restricted interval without an efficient envy-free allocation may select an envious outcome. The result applies to restrictions of this *one parent curve*; another economy may reveal another peak.

5.2. Representing the selected point with welfare indices

Interval choices reveal society's preferred coordinate τ_H, but do not yet give an interpersonal welfare representation of that choice. We now ask whether the *two CCE pairs* can select the revealed coordinate while preserving the separate envy tests. The indices using person 2's path agree at a_H, and those using person 1's path agree at b_H. Between these cutoffs the two pairs give opposite orderings of the reported levels. A weighted mixture therefore has an equality somewhere between them. By choosing its weight, we can place the equality at the already observed peak τ_H. The original two pairs, not this mixture, continue to test envy.

First choose continuous strictly increasing *calibrations* $\theta_1, \theta_2 : [0,1] \to \mathbb{R}$. Each changes the reported units on one recipient's path, and the same calibration applies to both evaluators on that path. We fix them *before* using the observed peak; taking both to be the identity is allowed. For a weight $\alpha \in [0,1]$, mix the two

calibrated CCE measures for person i:

$$z_i^{H,\alpha}(x) = \alpha\,\theta_1\Big(\pi_i^{1,H}(x)\Big) + (1-\alpha)\,\theta_2\Big(\pi_i^{2,H}(x)\Big), \qquad i = 1,2, \quad x \in D_H. \tag{6}$$

Because both CCEs rise with person i's preferences, their positive mixture still represents those preferences. But equality of the *mixed* levels is not the same as the two separate no-envy tests. It identifies a candidate for society's selected allocation. The directed comparisons remain those in (1) and (2), unless the two envy boundaries happen to coincide.

For any fair peak τ_H, the next result finds a weight that makes the two mixed levels meet there. The calibrations are already fixed; only the weight is chosen to fit the observed selection.

Proposition 5 (The two envy representations can express any selected fair point)**.** *Let H have an ordered contract curve, satisfy Assumption 1 on D_H, and have $a_H \le b_H$. Fix any θ_1, θ_2 as above. Each $z^{H,\alpha}$ represents preferences and has a unique equality at $t_\alpha^H \in [a_H, b_H]$. The map $\alpha \mapsto t_\alpha^H$ is continuous and onto, strictly increasing when $a_H < b_H$, with $t_0^H = a_H$, $t_1^H = b_H$; if $a_H = b_H$, all weights yield that point.*

With weight $\alpha = 0$, the mixture uses only person 2's reference pair and its two levels meet at a_H. With $\alpha = 1$, it uses person 1's reference pair and they meet at b_H. As α rises, the equality moves continuously through the fair interval. Thus, once the calibrations have been chosen, each peak in a nondegenerate interval determines one weight. For example, $\theta_k(r) = \widetilde{U}_k(C_k^H(r))$ expresses positions using the recipient's utility representation, while identity calibrations keep the original positions. Changing these scales changes the numerical mixture weight without changing the selected allocation. The weight is consequently a representation parameter, not an independently identified social welfare weight.

Suppose interval choices reveal $\tau_0 = 0.43$. Under identity calibrations, person 1's report exceeds person 2's by about 0.02113 using person 2's reference pair, but is lower by 0.01745 using person 1's pair. Giving the latter pair weight about 0.5477, and the former weight 0.4523, cancels these two differences. The mixed reports then meet at about 0.50288. The revealed social choice supplies the desired coordinate; the directed CCE pairs supply the two economically meaningful comparisons that are mixed.

Making two indices equal at τ_H does not by itself make a welfare aggregator select it. For example, adding the indices may favor a different coordinate after their units change, even if the change preserves individual rankings and their equality at τ_H. The reason is that the two slopes can change independently. The next calculation shows why the *minimum* of the two levels still selects their equality.

Consider two linear indices equal to c at $t = \tau$, one rising and the other falling as t increases. Their initial slopes could be 1 and -1. Change the units independently while preserving their common value c, so the slopes become $\lambda_1 > 0$ and $-\lambda_2 < 0$. Their *sum* now increases or decreases according to the sign of $\lambda_1 - \lambda_2$, possibly selecting an endpoint rather than τ. Their *minimum* behaves differently: the lower level rises as we approach equality from the left and falls after equality on the right. It therefore reaches its unique maximum at τ, whatever the positive slopes.

We ask which continuous monotone aggregator has this property for *every* pair of opposing continuous indices that meet once in the interior, including after independent slope changes that preserve the common equality. The proposition proves that such robust selection characterizes maximin up to an increasing transformation. It does not authorize arbitrary rescalings that move the indices' equality point [d'Aspremont and Gevers, 1977, Hammond, 1976].

Proposition 6 (Robust equalization characterizes maximin)**.** *Let $W : \mathbb{R}^2 \to \mathbb{R}$ be continuous, coordinatewise nondecreasing, and strictly increasing on the diagonal. These are equivalent.*

(i) *On every compact interval where continuous indices, one strictly increasing and one strictly decreasing, meet in the interior, W uniquely selects their equality.*

(ii) *For all $c, \tau \in \mathbb{R}$, $\lambda_1, \lambda_2 > 0$, and compact intervals containing τ in their interior, W uniquely selects τ for*

$$z_1(t) = c + \lambda_1(t - \tau), \qquad z_2(t) = c - \lambda_2(t - \tau).$$

(iii) *For some continuous strictly increasing ϕ,*

$$W(z_1, z_2) = \phi(\min\{z_1, z_2\}).$$

Why does the proposition select maximin? When one report rises and the other falls, only the lower report rises until the equality and falls afterward for every admissible pair of slopes. The theorem identifies all aggregators that preserve this behavior. It does *not* infer the desired equality from no-envy alone: observed interval choice first revealed the peak, and the mixture weight placed the two indices' equality there.

Theorem 4 (Maximin represents contraction-consistent envy-free choice)**.** *Fix H, $\mathcal{G}_H$, and the preceding calibrations. For a feasible choice correspondence C, these are equivalent.*

(i) *Choice is single-valued, Pareto-set dependent, contraction consistent, satisfies the specified one-sided continuity, and selects efficient no-envy when available.*

(ii) *Some $\alpha_H \in [0,1]$ makes the CCE mixture in (6) represent choice on every $G \in \mathcal{G}_H$:*

$$C(G) = \arg \max_{g \in \mathrm{PE}(G)} \min\{z_1^{H,\alpha_H}(g_1), z_2^{H,\alpha_H}(g_2)\}.$$

The unique choice is the projection in (5) at $\tau_H = t^H_{\alpha_H}$.

The interval-choice representation follows in three steps. Contraction consistency and one-sided continuity first reveal a peak τ_H, which envy-free selection on the full curve places inside $[a_H, b_H]$. After the CCE calibrations are fixed, a mixture weight places its equality at that peak; maximin then selects the equality when feasible and the nearest endpoint otherwise, reproducing *all interval restrictions* of the original curve. Merely shifting arbitrary own-utility reports might fit one selected allocation without ensuring that their equalities lie in the fair interval. The weight is unique for fixed calibrations when $a_H < b_H$; if $a_H = b_H$, every weight gives the same equality. None of this makes maximin the source of society's peak: the choice observations supplied it.

5.3. Choice on compact menus

The interval theorem does not settle choices when policy is restricted to a few *disconnected options*. Suppose an implementation constraint leaves only the cash–quality shares 0.425 and 0.433, both within the fair interval. Knowing that society chooses 0.43 from intervals does not determine which of these two alternatives it prefers. More seriously, a CCE mixture that correctly selects 0.43 from every interval can prefer an *envious* option to an available fair one on a two-point menu. In the running example, take identity calibrations, mixture weight $\alpha_H \approx 0.5477$, and write $m_\alpha(t) = \min_i z_i^{H,\alpha}(\gamma_{H,i}(t))$. On the menu $\{a_0, 0.439\}$, its scores are

$$m_\alpha(a_0) \approx 0.492260 \quad < \quad m_\alpha(0.439) \approx 0.493557.$$

Thus the mixture chooses $0.439 > b_0$, where person 2 envies, instead of the available envy-free boundary a_0. The interval theorem has not failed: this two-point menu is not an interval. We must impose new restrictions if one mixture is to represent *all* observed menu choices and remain fair whenever a fair alternative is offered.

Start with the two fair boundaries. If $m_\alpha(b_H) > m_\alpha(a_H)$, continuity makes an envious point just beyond b_H score higher than the fair point a_H. A two-point menu containing those alternatives would choose unfairly. The reverse inequality causes the analogous problem near a_H. Therefore a mixture that selects fairly from *every compact menu* must satisfy $m_\alpha(a_H) = m_\alpha(b_H)$. With this equality, fair interior positions score above envious exterior ones. But matching the boundary scores does not by itself reproduce all observed ties and preferences between positions on opposite sides of the peak.

We continue to study menus drawn from *one fixed efficient curve*, where deleting some efficient allocations cannot make a formerly dominated allocation efficient. On unrestricted feasible sets, by contrast, deleting alternatives may change which allocations are efficient. Sakamoto [2013] shows that efficient envy-free selection conflicts with contraction consistency on sufficiently rich finite domains. Our fixed-curve characterization does not claim to resolve that broader impossibility.

A two-point menu reveals whether society prefers one coordinate to another, even if neither is its peak. Let $c_h(K)$ be the chosen coordinates from any nonempty compact set $K \subseteq [0,1]$, where h labels the *parent allocation problem*. The set may be an interval, a singleton, or a disconnected menu such as $\{a_0, 0.439\}$. We may consider a fixed parent or a compact metric family P of parents whose preferences, recipient paths, and CCEs vary continuously and satisfy Range. Observing all such menus reveals the ties and comparisons across the peak that interval restrictions leave unresolved.

For the choices to reflect one social ranking, deleting an unchosen option should not overturn the choice, and a tie should not disappear simply because other options are removed. Sen's conditions α and β impose these two requirements. Together, when at least one previously chosen alternative remains in a smaller menu $L \subseteq K$, the chosen alternatives in L are exactly the surviving earlier choices:

$$c_h(L) = c_h(K) \cap L \quad \text{if } L \subseteq K, \qquad c_h(K) \cap L \neq \varnothing.$$

Condition α protects an alternative that survives a menu reduction; β protects all alternatives tied with it [Sen, 1971]. We additionally require one choice from each interval, and a *closed binary-choice graph*: whenever one coordinate is weakly chosen over another along a converging sequence of pairs, the limiting pair must retain that weak comparison. These assumptions yield a continuous single-peaked social ordering and a peak that varies continuously across parent problems. They do not yet show that a *CCE mixture* can numerically represent every comparison in that ordering.

Fair selection now constrains not only the peak but also comparisons between *two different fair alternatives*. Suppose $a_h < b_h$ and society strictly prefers b_h to a_h. By continuity, a slightly larger coordinate where person 2 envies would also beat a_h. Offered just those two points, society would select the envious one although the fair boundary remained available. Reversing the preference between the boundaries produces the analogous failure on the left. Thus the two fair boundaries must be *tied* and the peak must lie strictly between them when they differ. If $a_h = b_h$, that single fair point is the peak and many numerical weights can represent it.

Sometimes the two envy cutoffs are reversed, $a_h > b_h$. Then every efficient allocation makes at least one recipient envious, so the phrase "select an envy-free allocation when available" gives no guidance on these menus. We use a separate criterion, *strict balance*: either both recipients strictly prefer their own packages, both strictly prefer the other's packages, or both are indifferent. This strengthens Daniel's two-person balanced-envy requirement by excluding unilateral indifference [Daniel, 1975, Thomson, 2011]. On an ordered curve, the strictly balanced positions lie strictly between distinct cutoffs, or at their common point. With reversed cutoffs, these positions involve mutual strict envy. Strict balance therefore restricts choice when efficient no-envy is unavailable, without claiming that mutual envy is the same as no-envy.

Even choices satisfying these fairness requirements need not have the *particular CCE representation* we want. A mixed CCE score is built from only two increasing calibration functions. To the left of its peak, person 1's mixed level is lower; to the right, person 2's is lower. If society ties a position on the left with one on the right, the lower score on each side must be equal. Many such ties and strict comparisons therefore constrain the *same two functions*; fitting each observed comparison with different calibrations would not represent one social ordering.

The consistency problem resembles cancellation in additive measurement. Suppose three observed comparisons require changes in two calibrated coordinates of $(2,-3)$, $(2,-1)$, and $(-4,4)$, and each comparison is said to increase the same additive score. Adding the three changes gives $(0,0)$. Their supposed three strict score increases would therefore sum to a positive change in a score whose *total change is zero*, an impossibility. The actual CCE theorem must test analogous collections of comparisons, including ones obtained only as limits of longer collections. The Online Appendix states the required *closed-cone cancellation* condition in (C.5) and relates it to additive conjoint measurement [Luce and Tukey, 1964]. The Additional Material gives an example in which finite checks alone miss the obstruction.

Theorem B (CCE rationalization of compact-menu choice)**.** *Under the standing ordered-path, continuity, and Range conditions, the following are equivalent.*

(i) *Choice satisfies Sen's α and β, single-valued interval choice, a closed binary-choice graph, and closed-cone cancellation. It selects envy-free allocations when $a_h \leq b_h$ and available, or strictly balanced allocations when $a_h > b_h$ and available.*

(ii) *Every continuous $\alpha : P \to (0,1)$ admits jointly continuous, strictly increasing calibrations normalized at zero and one, with*

$$z_i^h(x) = \alpha_h \theta_1(h, \pi_i^{1,h}(x)) + (1-\alpha_h)\theta_2(h, \pi_i^{2,h}(x)).$$

The minimum $m_h(t) = \min_i z_i^h(\gamma_{h,i}(t))$ represents all compact-menu choices and ties, with

$$c_h(K) = \arg\max_{t \in K} m_h(t) \quad (K \neq \varnothing \text{ compact}), \qquad m_h(a_h) = m_h(b_h).$$

Theorem B separates consistency of observed choices from their representation by CCEs. Sen's conditions, unique interval choice, and the closed binary-choice graph provide a continuous social ordering, including its ties; envy-free or strict-balance selection restricts which alternatives may outrank an eligible one. *Closed-cone cancellation* further requires that one pair of increasing CCE calibrations satisfy *all* revealed comparisons simultaneously, including their limiting implications. When it can, maximin of the calibrated mixture represents every compact-menu choice, with the two boundary scores equal. Conversely, the existence of these calibrations for even one continuous interior coefficient family yields the stated conditions. This conclusion cannot be inferred from interval choices alone.

Notice the order in which the numbers are chosen. For interval menus, we fix the calibrations first and choose a weight to reproduce the revealed peak. For all compact menus, we observe the *whole social ordering* first, prescribe *any continuous interior coefficient family*, and construct calibrations that reproduce its ties and comparisons. The freedom to choose the coefficient does not identify an underlying social welfare weight; it shows that the same ordinal choices have multiple equivalent numerical representations.

5.4. Joint rationalizability across economic problems

Even after the compact-menu problem is solved within each economy, separate numerical CCE normalizations need not fit into *one comparison of policies across economies*. The next characterization asks when all the local representations can be reconciled without changing their preference rankings or the fairness tests. Consider the compulsory benefit replacement studied below. The original allocation gives both recipients $(0.5, 0.5)$; the replacement gives $(0.3, 0.6)$ to person 1 and $(0.4, 0.6)$ to person 2. Own-utility reports remain $(1.5, 1)$, but person 1 begins to envy. If the planner is required to rank the allocations differently because fairness changed, yet is also required to treat identical own-utility reports as identical social alternatives, no ranking can satisfy both demands. The conflict depends on *what the planner is allowed to observe*.

We therefore specify the information rather than assume it. Let $Q(h, x)$ be the report available about alternative x in parent problem h, and let Y collect the alternatives the planner might compare. An information requirement can force two descriptions to be tied when their reports are indistinguishable. But a score $V(h, Q(h, x))$ still observes the problem label h, so identical report values in different parents need *not* receive equal welfare ranks unless the requirement explicitly says so. A score depending only on Q would impose an additional cross-problem equality.

A common social ranking must respect the comparisons revealed by choices within each parent, any additional comparisons required by a welfare criterion, and ties imposed by the planner's information. For example, the welfare criterion may favor an efficient improvement, while two descriptions that cannot be distinguished from the available reports may have to receive the same rank. These instructions can contradict one another after several steps: if $x \succsim y$, $y \succsim z$, and $z \succ x$ are all required, transitivity yields $x \succsim z$, contradicting the required strict preference for z. Even if a contradictory comparison arises only from *limits* of observed comparisons, continuity can create the same difficulty. We therefore collect the forced weak comparisons, their transitive consequences, and their limits in the smallest *closed preorder* $\mathcal{R}$. Let $\mathcal{T}$ denote the strict comparisons the planner must respect. A common ranking is impossible if $\mathcal{R}$ includes the reverse of a comparison in $\mathcal{T}$.

Theorem C (Joint rationalizability across economic problems)**.** *Maintain the information requirements and local conditions of Theorem B, except cancellation. These are equivalent.*

(i) *Jointly continuous local CCE representations and a continuous $W : Y \to \mathbb{R}$ preserve local choices, ties, required comparisons, and report restrictions.*

(ii) *Closed-cone cancellation holds, and the closed preorder reverses no required strict comparison:*

$$(y, x) \notin \mathcal{R} \qquad \text{for every } (x, y) \in \mathcal{T}.$$

On the attained domains,

$$W(x) = V(h, Q(h, x)), \qquad W(\gamma_h(t)) = \rho(h, m_h(t)),$$

for continuous V, ρ with $\rho(h, \cdot)$ strictly increasing.

Theorem C requires consistency at two levels. Within each economy, closed-cone cancellation asks whether one CCE mixture represents all observed menu choices. Across economies, the no-reversal condition asks

whether the resulting local rankings can coexist with required welfare comparisons and ties imposed by the reports. Checking direct pairs alone is not enough, because a sequence of transitive comparisons or a limiting sequence may create the contradiction. The proof uses a continuous representation theorem for closed preorders [Evren and Ok, 2011, Theorem 1 and Corollary 2]. Under the stated conditions, any continuous interior coefficient family is admissible, and maximizing the common score on compact menus yields choices with Sen's α, β, and upper hemicontinuity. Equal local weights neither ensure nor are needed for cross-problem consistency.

The theorem does not automatically enforce Pareto efficiency or fair choice on *new menus beyond the original curves*. Such requirements must be included among the comparisons the global ranking is asked to satisfy. For instance, a strict Pareto improvement must beat a dominated allocation, and an eligible alternative must outrank an ineligible one in the relevant comparison if eligible choice is required. Here "eligible" means satisfying the criterion chosen for that environment, such as envy-freeness or strict balance. The common ranking then respects those additional requirements when they are consistent; away from the original curves it need not equal the raw maximin CCE score. On each original curve, the increasing transformation ρ preserves the local ordering.

Finally, the factorization $W(x) = V(h, Q(h, x))$ keeps the parent label available. Requiring welfare to depend on Q *alone* would force more alternatives from different parents to tie. That stronger information restriction can be economically consequential when a reform changes recipients' packages without changing their own-utility reports. We next study exactly how such changes move the envy cutoffs and the fair allocation set.

6. STRUCTURAL CHANGE AND PUBLIC-BENEFIT REFORM

The characterization has an immediate implication for counterfactual welfare analysis: changing the physical packages recipients can receive may change the references required by no-envy, even when preferences remain fixed. An old normalization then need not report the new directed comparisons correctly. We first establish when an inherited index remains valid and how improving one recipient's packages shifts both envy cutoffs. We next use a linear cash–service reform, distinct from the nonlinear cash–quality example used for the basic construction, to isolate the missing information. The reform preserves the *entire attainable own-utility possibility set* and public spending, yet can move or eliminate the efficient envy-free set. This establishes what an analyst restricted to own-welfare possibilities may overlook.

6.1. When an inherited welfare comparison remains valid

If a policy merely removes points from the original contract curve, the physical packages at the surviving points are unchanged. The existing recipient-based CCEs therefore continue to express their envy comparisons, and the earlier interval-choice representation still applies. A reform that *replaces the package delivered at a coordinate* is different. Changes in technology, service composition, or delivery probabilities can make person 2's new package worth something else to person 1. An index may still rank each person's own packages correctly yet fail to give the new recipient packages common interpersonal levels.

When can the old indices still express a complete directed envy comparison? Suppose C_k^0 is recipient k's old efficient assignment path and C_k^1 her new one. The same continuous indices e_1, e_2 are used in both economies, and the two evaluators agreed on the level of every old recipient package. To agree on the *new* path as well, they need a shared indifferent match: for each new package there must be one old package that *both people* regard as equally desirable. The following condition expresses this match through an increasing map ρ. If

$$e_1(C_k^0(r)) = e_2(C_k^0(r)) \qquad (r \in [0,1])$$

and there is a continuous strictly increasing map ρ satisfying

$$C_k^1(r) \sim_i C_k^0(\rho(r)) \qquad (i = 1, 2)$$

throughout the counterfactual support, then the inherited indices also agree on C_k^1. Conversely, if the same indices agree on both paths and use the same common range of welfare levels, such a common increasing map exists.

The map ρ identifies, for every new recipient package, an old package that both evaluators regard as equally good. Since the indices gave the old package a common level and preserve preferences, they must give the new package that same level. If each evaluator instead matches it to a *different* old position, common levels need not survive. The converse needs equal *complete ranges* of common welfare levels; agreement over only a short portion of the old and new paths is insufficient. Unchanged individual preference rankings alone therefore do not certify an inherited interpersonal measure.

Now suppose person 2 receives a better package at every coordinate according to *both* evaluators, while person 1's packages remain the same. To stop envying, person 1 needs a higher coordinate because person 2's package has improved. Person 2 can also accept a higher coordinate for person 1 before beginning to envy, since she now receives a better package herself. Thus both directed cutoff positions move to the right. The next proposition states the prediction and the exact conditions for strict movement.

Proposition 7 (Improving one recipient's assignments shifts the no-envy interval)**.** *For $s = 0, 1$, let $\gamma^s(t) = (C_1(t), C_2^s(1-t))$ be ordered efficient curves with cutoffs a_s, b_s. If C_1 is unchanged and $C_2^1(r) \succsim_i C_2^0(r)$ for both evaluators and all r, then*

$$a_1 \geq a_0, \qquad b_1 \geq b_0.$$

The first inequality is strict if $C_2^1(1-a_0) \succ_1 C_2^0(1-a_0)$; the second is strict if $C_2^1(1-b_0) \succ_2 C_2^0(1-b_0)$.

The proposition compares the packages delivered at the *same coordinate*. Person 1 still receives $C_1(t)$, but person 2 now receives $C_2^1(1-t)$ instead of $C_2^0(1-t)$. If the old selected coordinate τ_0 was fair, the new fair set is nonempty, and $\tau_0 < a_1$, person 1 envies at the old coordinate after the reform. Every new efficient envy-free allocation must give person 1 a coordinate at least $a_1 > \tau_0$. A planner cannot safely carry the old choice into the new problem just because person 1's own package has not changed.

We can see the same effect in CCE numbers. Fix some package x. If every package on person 2's new path is better than the old package at the same coordinate, an evaluator reaches the utility of x at a *lower* position on the new path. Thus $\pi_i^{2,1}(x) \leq \pi_i^{2,0}(x)$. Relative to person 2's actual assignment, person 1's unchanged package now has a lower equivalent, pushing the point where she stops envying to the right. Person 2's improved package also lets her tolerate more resources for person 1 before she envies, moving the other cutoff right.

The movement can even reverse the fairness of two otherwise comparable policies. Suppose each environment has exactly one efficient envy-free coordinate, t_0 before and t_1 after the improvement. If person 1 strictly prefers the improved package person 2 gets at t_0, then she envies there, so the new equality has $t_1 > t_0$. At t_1 in the old economy, person 2 envies. Each environment therefore favors a different allocation under the *same* no-envy principle. The reversal arises from changed recipient packages, not from changing tastes.

Return to cash and service. Keep $s + r \leq 1$ and $U_i(m, q) = m + \beta_i q$, where $\beta_1 > \beta_2 > 0$. Person 1's path remains $C_1(r) = (r, r)$. Change only person 2's service technology, yielding the path:

$$C_2^\kappa(r) = \big(r, \kappa r + (1-\kappa)r^2\big), \qquad \kappa \in [0, 1].$$

A larger κ adds service to person 2's package at each interior resource share; it changes neither person 1's package nor either person's preferences.

For $(\beta_1, \beta_2) = (2, 1)$, consider person 1's share $t = 0.43$. Her package $(0.43, 0.43)$ is worth 1.29 to her. With $\kappa = 0$, person 2's package $(0.57, 0.3249)$ is worth 1.2198 to person 1, so she does not envy. With $\kappa = 1/2$, person 2 instead receives $(0.57, 0.44745)$, which person 1 values at 1.4649. She now envies despite receiving the same package. Person 2's own welfare also rises in this example; the next reform will change package composition while holding *each recipient's own utility constant*.

The two cutoffs rise for different reasons. Person 1 needs more input to match person 2's improved package; person 2, receiving the improvement, is willing to let person 1 have more input. With $\beta_1 > \beta_2$, both boundaries rise as κ increases, remain distinct for $\kappa < 1$, and meet at equal division when $\kappa = 1$. The proof calculates the movements directly.

Both cutoffs moving right does *not* automatically make the old and new fair intervals disjoint. For disjointness, the new lower cutoff must exceed the old upper cutoff. Nor does granting person 1 a higher resource share necessarily hurt person 2, because person 2's new technology delivers more service from her remaining input. At full access, person 1 prefers the new equal split to any old fair allocation; person 2 does so exactly when $\beta_1(1-\beta_2) < 2\beta_2$. Under this additional condition the new fair allocation Pareto dominates the

old fair set. These statements concern the value of delivered packages, not the labels assigned to resource shares.

For $(\beta_1, \beta_2) = (2, 1)$, the initial fair interval is $[0.4189, 0.4384]$. At $\kappa = 1/2$, it becomes $[0.4586, 0.4689]$; at full access it is the equal split. The intervals become disjoint once $\kappa > 1/4$. At full access both people prefer the new fair allocation to every old fair allocation. Thus improved opportunities can shift the fairness comparison even when the original recipients' tastes are unchanged.

A change in delivery risk can create the same kind of problem. In the contingent-payment economy, suppose the planner calibrated an index pair when delivery probability was p_0, then kept that pair when probability became p. With reference lotteries extended to all nonnegative promised payments, maximin of the *old* pair now gives person 1 the share $t^{\text{old}}(p) = p/(p + p_0)$, although both people still agree that a fair split is one half. As p and p_0 vary, this inherited recommendation can lie anywhere in the share interval. Without the extended reference path, the equivalents may even be undefined for $p > p_0$. The example does not preserve own-utility opportunities; it isolates the danger of keeping a previously valid interpersonal reference when the physical contingent contracts change.

6.2. Cash and in-kind public benefits

The service improvement above changes person 2's welfare as well as the package person 1 compares with her own. Can envy change when *neither person's own welfare changes*? Reference assignments already matter in the theory of public and nonrival goods: Maniquet and Sprumont [2004] derive a welfare measure based on the quantity of an excludable nonrival good an individual would accept at no cost. Our setting instead uses recipient-specific cash–service packages to evaluate no-envy. Consider a program that compulsorily replaces cash with additional in-kind service. The program deducts just enough cash to leave *each recipient* indifferent to her old package, and it keeps total spending fixed. Another person need not agree about the value of that replacement. The resulting difference in *cross-package values* isolates the fairness effect of changing benefit composition.

Let m denote cash and q service. Person i's preferences are $U_i(m, q) = m + \beta_i q$, with person 1 valuing service more, $0 < \beta_2 < \beta_1$. Delivering a package costs the program $c_p(m, q) = m + pq$, where one unit of service costs p and $\beta_2 < p < \beta_1$. Replacing cash with service while keeping person 1 indifferent saves public money because she values service above its cost; doing the same for person 2 costs more because she values service below its cost. We use these opposite budget effects to preserve total expenditure while making both people indifferent to their own replacements.

Begin with the equal division $r = 0.5$, giving each person cash and service $(0.5, 0.5)$. Set $(\beta_1, \beta_2) = (2, 1)$ and service cost $p = 1.5$. Give each person 0.1 extra unit of service, but remove from her cash exactly what that service is worth *to her*: 0.2 from person 1 and 0.1 from person 2. Denote the replacement of recipient i's package at coordinate r by $C_i^\delta(r)$. At the equal split the new packages are

$$C_1^\delta(0.5) = (0.3, 0.6), \qquad C_2^\delta(0.5) = (0.4, 0.6).$$

Check four valuations and the public bill. Person 1's own utility remains $0.3 + 2(0.6) = 1.5$, and person 2's remains $0.4 + 0.6 = 1$. The packages cost 1.2 and 1.3, so total spending is still 2.5. But person 1 values *person 2's package* at $0.4 + 2(0.6) = 1.6 > 1.5$ and envies, whereas person 2 values *person 1's package* at $0.3 + 0.6 = 0.9 < 1$ and does not. Each person was compensated for the service using her *own* valuation; the other person's valuation of that same service is different. The missing cross-evaluations, not the own utilities, explain the new envy.

The effect extends beyond the two packages at the equal split. Before replacement, equal division is the unique efficient envy-free allocation; after replacement, the efficient fair shares lie around $[0.5166, 0.5249]$. The two envy boundaries have separated, so two directed pairs of indices are now needed on the efficient curve. To show that this is more than a change at one selected allocation, we next construct replacements *at every resource coordinate* and verify that the entire own-utility possibility set remains the same.

Before reform, recipient i's package at resource coordinate r is $C_i^0(r) = (r, r)$. Let $v(r)$ describe the extra service: it is continuous, positive for interior shares, symmetric around one half, and zero at the endpoints. Recipient i receives $\delta_i v(r)$ additional service. To leave her indifferent, the program removes $\beta_i \delta_i v(r)$ from

cash. The resulting compensated path is:

$$C_i^{\delta}(r) = \big(r - \beta_i \delta_i v(r),\, r + \delta_i v(r)\big).$$

The compensation works at *every* r, not just at $1/2$. Because each person values an extra unit of service at β_i, the cash deduction cancels its utility exactly:

$$U_i(C_i^{\delta}(r)) = U_i(C_i^{0}(r)) = (1 + \beta_i)r.$$

Individual indifference is not yet fiscal balance. Replacing person 1's cash with service saves public money because $\beta_1 > p$; replacing person 2's costs money because $\beta_2 < p$. To cancel these effects, choose the following ratio of service adjustments:

$$\delta_2 = \frac{\beta_1 - p}{p - \beta_2}\delta_1.$$

With this ratio, total spending along the efficient curve remains unchanged. The magnitudes must be sufficiently small to keep cash and service feasible, preserve the common ordering of both recipient paths, and make public spending increase with the resource coordinate. These conditions protect the intended feasible technology; they do not assume away differences in cross-package valuations.

Now compare the new packages *across* recipients. Person 1 values person 2's additional service more than person 2 does, creating a positive cross-package gain $(\beta_1 - \beta_2)\delta_2 v(r)$ that makes person 1 demand a larger resource share before she stops envying. Person 2 values person 1's added service less than person 1 does, creating the corresponding cross-package loss $(\beta_1 - \beta_2)\delta_1 v(r)$; she can tolerate a larger share for person 1 before she starts envying. The budget condition fixes the ratio δ_2/δ_1, which depends on the service price. When person 1's lower cutoff exceeds person 2's upper cutoff, no efficient allocation passes both tests.

To find the exact cutoff coordinates, the two envy equalities compare the baseline resource difference $2r - 1$ with a service adjustment of size $v(r)$. For $r \in (1/2, 1)$, define their ratio

$$\psi_v(r) = \frac{2r - 1}{v(r)} \qquad (r \in (1/2, 1)).$$

Assume ψ_v rises strictly from zero to infinity as r moves from one half to one. Its inverse Ψ_v then turns each required envy-gap ratio into a unique cutoff coordinate. We also want to compare *every feasible own-utility pair*, not only those on the efficient curve. For $E = 0$ before the replacement and $E = \delta$ after it, define the utility possibility set by taking all package pairs whose *total public cost* is within the budget:

$$\mathcal{U}^E = \left\{\big(U_1(C_1^E(s)), U_2(C_2^E(t))\big) : c_p(C_1^E(s)) + c_p(C_2^E(t)) \le 1 + p\right\}.$$

Theorem 5 (Benefit composition changes the envy-free set while preserving the utility possibility set)**.** *Under the preceding conditions and budget* $1 + p$*, both environments satisfy*

$$c_p(C_1^E(s)) + c_p(C_2^E(t)) \le 1 + p \quad \Longleftrightarrow \quad s + t \le 1$$

and $\mathcal{U}^0 = \mathcal{U}^{\delta}$*. Efficient allocations exhaust the budget; initially the only efficient envy-free share is* $s = 1/2$*.*

Write $\Delta = \beta_1 - \beta_2$ *and*

$$p^* = \frac{\beta_1 + \beta_2 + 2\beta_1\beta_2}{2 + \beta_1 + \beta_2}.$$

After replacement the envy cutoffs are

$$a = \Psi_v\left(\frac{\Delta\delta_2}{1 + \beta_1}\right), \qquad b = \Psi_v\left(\frac{\Delta\delta_1}{1 + \beta_2}\right).$$

The efficient envy-free set is empty for $p < p^*$*, the singleton* $\{a\} = \{b\}$ *for* $p = p^*$*, and the nondegenerate interval* $[a, b]$ *for* $p > p^*$*.*

The critical price p^* tells us whether the two people can both be free of envy on an efficient allocation. Below p^*, the resources person 1 needs to stop envying exceed what person 2 will accept, so the efficient fair set is empty. At p^*, the cutoffs coincide and exactly one efficient allocation is fair; above it, they bound a nondegenerate fair interval. For $(\beta_1, \beta_2) = (2, 1)$, $p^* = 1.4$, placing our example at $p = 1.5$ in the interval case. Only at the price where the cutoffs coincide can one index pair represent both directed tests *on the efficient curve*. The threshold reflects the recipients' different valuations of service and the balancing of their compensation costs, rather than the size of an admissible adjustment. The Online Appendix gives explicit formulas.

The theorem is not just about envy at one package pair. At every feasible resource coordinate, each recipient receives *exactly the same own utility* before and after replacement. The budget feasibility condition is likewise unchanged. Therefore *every fixed preference-representing index* assigns each old and new recipient package the same level, so all attainable pairs of own-welfare reports are identical. To state this information result for arbitrary fixed preference representations e_1, e_2, write

$$\mathcal{W}_e^E = \left\{ \left(e_1(C_1^E(s)), e_2(C_2^E(t)) \right) : s + t \leq 1 \right\}.$$

The sets $\mathcal{W}_e^0$ and $\mathcal{W}_e^\delta$ are identical: a planner who sees only the complete set of possible own-welfare reports cannot tell whether the packages have been replaced. The planner nevertheless needs different cross-package comparisons to find the efficient envy-free set.

Corollary 1 (The utility possibility set does not determine efficient no-envy)**.** *In Theorem 5, preferences and spending are unchanged, and $\mathcal{W}_e^0 = \mathcal{W}_e^\delta$ for every fixed preference-representing pair e, yet the efficient envy-free sets differ as in Theorem 5. Thus no single-valued rule based only on these own-welfare possibility sets implements efficient no-envy in both environments.*

The economic implication depends on the price regime. For $p > p^*$, both economies have nonempty but disjoint efficient envy-free sets, despite identical own-welfare reports. A rule using only those reports cannot implement a fair choice in each environment. For $p < p^*$, the new economy has *no* efficient envy-free allocation. That failure is not a special information problem: even a fully informed planner cannot select an allocation that does not exist. At $p = p^*$ a single fair coordinate remains. Thus the reform separates the information lost in own-welfare measurement from the outright incompatibility of the two directed fairness requirements.

In the linear cash–service model, the planner can recover both missing envy comparisons without separately asking for all four utilities. For a budget-exhausting allocation, suppose she knows the public budget B, the service price p, both coefficients β_i, the recipients' own-utility reports $v_i = U_i(g_i)$, and total service $Q = q_1 + q_2$. Because $m_1 + m_2 = B - pQ$, the difference between the value of a person's own package and her value of the other's is

$$U_i(g_i) - U_i(g_j) = 2v_i - B + (p - \beta_i)Q, \qquad j \neq i.$$

This formula shows exactly what extra physical information is useful: given own utilities and Q, the planner can calculate the two envy gaps. Because $v_1 + v_2 = B + (\beta_2 - p)Q + (\beta_1 - \beta_2)q_1$, and $\beta_1 \neq \beta_2$, she can also recover q_1, then $q_2 = Q - q_1$ and $m_i = v_i - \beta_i q_i$. Total service is therefore a sufficient extra statistic in *this linear model*; it need not be sufficient when valuations are nonlinear.

The result is not tied to both recipients receiving extra service according to an identical adjustment shape. The Online Appendix allows a general positive schedule for person 1 and determines person 2's schedule from reflection and fiscal balance. The efficient envy-free set remains nonempty exactly when $p \geq p^*$, subject to the stated ordering and cost conditions. Above the threshold it contains an open interval; at the threshold uniqueness needs an additional monotonicity restriction. Unrelated adjustment schedules need not give the same boundary geometry.

Exact compensation makes the missing cross-package valuations particularly clear. The same distinction survives small errors in compensation, as the next extension shows. Suppose person 1 loses an additional $\lambda v(r)$ in cash, while person 2 receives the opposite cash adjustment. The public budget and feasible resource shares are unchanged, but the own-utility possibility set moves by at most $|\lambda| \max_r v(r)$ in Hausdorff distance, meaning every point of one set lies within that distance of some point of the other. For sufficiently small errors, the strict results remain: at $p > p^*$, the new fair interval is still nonempty and disjoint from the old one;

at $p < p^*$, efficient no-envy remains impossible. The Online Appendix proves these claims and derives the perturbed cutoffs.

The replacement has so far been *compulsory*. If recipients may keep their old packages, the feasible set contains more alternatives and may have a different efficient frontier. In the next corollary, retaining the old options permits an efficient envy-free allocation that both recipients strictly prefer to the initial equal division. Thus the disappearance result cannot be read as a claim that a voluntary improvement necessarily destroys efficient fairness.

Corollary 2 (Keeping old benefit options restores efficient no-envy)**.** *Let the planner retain baseline and reformed packages for both recipients, subject to the same public budget. Put $A = (\beta_1 - p)\delta_1$ and $P = 1 + p$. An efficient envy-free allocation gives person 1 her reformed package and person 2 his baseline package at*

$$r = \Psi_v(A/P) > \frac{1}{2}.$$

Both strictly prefer it to the original fair allocation.

For an agency applying no-envy, the model shows why unchanged own-welfare reports need not settle whether a reform is fair; it does not estimate any program's effect. Cash-out experiments and WIC food-package revisions change the form or composition of benefits, while public-housing assignments allocate heterogeneous in-kind packages [U.S. Department of Agriculture, Food and Nutrition Service, 1993, 2024, Waldinger, 2021]; related studies assess welfare and targeting effects [Currie and Gahvari, 2008, Gasparini and Pinto, 2006]. An agency adopting no-envy should reassess how recipients value each other's packages after a change, even when own welfare and spending stay fixed, and distinguish compulsory replacement from retained options. None of these studies shows that beneficiaries actually use our two no-envy comparisons.

One way to study the mechanism empirically would be to change assigned cash-service packages while keeping recipients' own valuations and program spending approximately constant. Elicit each person's preference between her own package and the other's, and separately ask which allocations seem acceptable. The one-sided improvement predicts the direction of both envy-cutoff movements; the compensated replacement predicts a change even when own utility is fixed. Such observations could distinguish this mechanism from an account using only stable own-utility reports. The studies cited above motivate the setting; they do not provide this test.

7. RISK SHARING

Risk sharing provides a second environment in which the normalization matters. A planner can vary a sure transfer along an efficient risk-sharing contract, while two recipients with different risk attitudes may rank the *same risky contract* differently. Equalizing their own certainty equivalents then need not satisfy no-envy. We calculate both cross-contract evaluations and recover the two directed tests in money units. In contrast, Fleurbaey and Zuber [2017, 2021] study fairness-motivated social evaluations and utility normalizations under different normative requirements. Our question is the narrower no-envy comparison of the actual risky contracts. The CARA model makes its fair-transfer interval explicit; the result does not say that any arbitrary insurance premium generates an ordered contract curve.

Proposition 8 (Equal own certainty equivalents can conceal envy)**.** *Suppose both people use the same belief and cash-normalized certainty equivalents. Person 2 is strictly more risk averse: $\mathrm{CE}_1(Y) > \mathrm{CE}_2(Y)$ for every nonconstant lottery Y. If the contracts satisfy $\mathrm{CE}_1(c_1) = \mathrm{CE}_2(c_2)$ and c_2 is nonconstant, then person 1 strictly prefers person 2's contract and envies.*

The proposition requires only one calculation on *person 2's risky contract*. Because person 2 is more risk averse, person 1 values that same nonconstant contract more highly: $\mathrm{CE}_1(c_2) > \mathrm{CE}_2(c_2)$. Equal own equivalents give $\mathrm{CE}_2(c_2) = \mathrm{CE}_1(c_1)$. Combining these inequalities yields $\mathrm{CE}_1(c_2) > \mathrm{CE}_1(c_1)$, so person 1 would exchange contracts and envies. Both own certainty equivalents were computed correctly; they simply did not include the required cross-contract valuation.

The example extends to efficient risk-sharing contracts. Under constant absolute risk aversion (CARA), efficient sharing divides the total risky resources in fixed proportions. A separate *sure cash transfer* increases

person 1's consumption and reduces person 2's by the same amount in every state. Person 1's envy comparison sets a minimum acceptable transfer; person 2's sets a maximum. Their intersection, if nonempty, is the interval of efficient envy-free transfers.

Let the two people share a common probability belief p with positive probability on each of finitely many states. Total consumption B varies with the state. Person i has CARA utility $u_i(c) = -e^{-\eta_i c}$, with $\eta_i > 0$ measuring absolute risk aversion, and assume the two coefficients differ. The efficient contracts split the risk in fixed proportions and add a state-independent transfer d:

$$c_1 = \alpha B + d\mathbf{1}, \qquad c_2 = (1-\alpha)B - d\mathbf{1}, \qquad \alpha = \frac{\eta_2}{\eta_1 + \eta_2}.$$

The risk-sharing fraction $\alpha = \eta_2/(\eta_1 + \eta_2)$ assigns person 1 the proportion of aggregate risk implied by the two risk tolerances [Wilson, 1968]. The sure transfer d moves the same cash amount to person 1 in each state, subtracting it from person 2. We need *four* certainty equivalents. Let $A_i = \mathrm{CE}_i(\alpha B)$ be evaluator i's value of person 1's risky share, and $D_i = \mathrm{CE}_i((1-\alpha)B)$ her value of person 2's. Cash additivity gives $\mathrm{CE}_i(c_1) = A_i + d$ and $\mathrm{CE}_i(c_2) = D_i - d$. Each evaluator can therefore compare both contracts in her own sure-money units.

With two equally likely states, take $B = (8, 12)$ and $(\eta_1, \eta_2) = (0.5, 1)$. If the planner equalizes only the *own* certainty equivalents $A_1 + d = D_2 - d$, she selects $d \approx -1.5629$. Person 1 then receives $(3.7704, 6.4371)$ and person 2 receives $(4.2296, 5.5629)$. Their own certainty equivalents are both about 4.6888. Yet person 1 values person 2's *same contract* at 4.7871, more than her own, and thus envies by 0.0984 in sure-money units.

To obtain a fair transfer, let each person compare *her own contract with the other's*. Person 1 needs $A_1 + d \geq D_1 - d$; person 2 needs $D_2 - d \geq A_2 + d$. Rearranging gives a lower bound on d from person 1 and an upper bound from person 2. The theorem identifies the resulting interval and shows that maximin of uncorrected *own* certainty equivalents selects outside it when risk attitudes differ.

Theorem 6 (Heterogeneous risk attitudes require directed certainty-equivalent comparisons)**.** *Under the specified common full-support beliefs, nonconstant aggregate B, and distinct positive CARA coefficients, efficient envy-free transfers are*

$$d \in [d_1, d_2], \qquad d_i = \frac{D_i - A_i}{2}, \qquad d_1 < d_2.$$

The directed no-envy comparisons are represented by

$$z^2 = \left(\mathrm{CE}_1(c_1),\ \mathrm{CE}_2(c_2) + D_1 - D_2\right),$$
$$z^1 = \left(\mathrm{CE}_1(c_1) + A_2 - A_1,\ \mathrm{CE}_2(c_2)\right).$$

Their equality points are d_1, d_2, respectively. Maximin of unadjusted own certainty equivalents instead selects $d^{\text{own}} = (D_2 - A_1)/2$, below d_1 when $\eta_1 < \eta_2$ and above d_2 when $\eta_1 > \eta_2$; thus the selected contract is envious.

The two corrections in the theorem supply the missing *cross-contract* evaluations. For person 1's envy, her own report is $A_1 + d$, while person 2's own report $D_2 - d$ must be adjusted by $D_1 - D_2$ to become *person 1's* valuation $D_1 - d$. For person 2's envy, the adjustment $A_2 - A_1$ converts person 1's own report into *person 2's* valuation of person 1's contract. The corrections therefore come from the four contract valuations, not from choosing a desired transfer. Mixtures of the two directed pairs express every fair transfer. In the numerical example, the fair interval is $[-1.5137, -1.4167]$, while own-CE maximin selects below it.

The requirement to use *money units* further restricts the indices. If giving a person one extra sure dollar in every state must increase her welfare report by exactly one, any preference-representing index has to equal her certainty equivalent plus a constant; the main proof appendix establishes this in Lemma 1. Requiring both evaluators to report the same level for every contract on the *compared recipient's cash-transfer path* then fixes the difference between these constants. A common additive constant has no effect. The directed corrections are thus fixed by the risk-sharing contracts and the common money metric.

The result assumes finite-state common beliefs, CARA preferences, and unrestricted sure transfers; a nonnegative-consumption requirement or transfer bounds can truncate the fair interval. Risk-sharing research has studied heterogeneous beliefs with exposure constraints [Boonen and Ghossoub, 2020] and the existence

of Pareto optima without convexity [Liebrich, 2024]. The Additional Material considers a different question: no-envy when recipients have identical risk attitudes but different subjective beliefs. Each person evaluates both risky contracts using *her own* probabilities. A common probability selected only to express the comparison is a *reference probability*, not a belief attributed to both recipients [Eden, 2020, Billot and Qu, 2021, Sprumont, 2019]. The distinction between a person's own equivalent and her valuation of the other contract remains essential.

8. CONCLUSION

The paper characterizes the equivalence-based welfare normalizations that represent envy freedom on ordered bilateral contract curves with Range. The relevant reference for a directed comparison is the efficient assignment path of the recipient being compared. Contract Curve Equivalents reproduce the two no-envy tests, generally through two distinct index pairs; a single pair works on the efficient curve exactly when the cutoffs coincide. This is the constructive answer to the limitations of fixed own-welfare reports under the specified information restriction. Relative-cost variation makes the numerical CCE form necessary, and additional invariance conditions identify the used physical references. Choosing among fair policies requires further social-choice restrictions; no-envy alone does not choose a particular point.

For a policy analyst who adopts no-envy, the results suggest a concrete evaluation procedure: verify the ordered-domain and coverage conditions, obtain each recipient's valuation of both actual packages, compute the two envy cutoffs, and only then apply a separate rule to choose among fair alternatives. Following a change in technology, benefits, or risk, those comparisons must be recomputed rather than inferred from unchanged own-welfare profiles. The compensated-reform theorem shows why the distinction matters even when all attainable own-utility profiles are unchanged; the treatment of optional old packages matters too. A follow-up empirical study could elicit cross-package preferences before and after a change in benefit composition and test the predicted cutoff movements. These recommendations depend on adopting no-envy, not on a claim that the theory determines one universally optimal tax or benefit policy.

REFERENCES

Alberto Alesina and George-Marios Angeletos. Fairness and redistribution. *American Economic Review*, 95(4):960–980, 2005. doi:10.1257/0002828054825655.

Antoine Billot and Xiangyu Qu. Utilitarian aggregation with heterogeneous beliefs. *American Economic Journal: Microeconomics*, 13(3):112–123, 2021. doi:10.1257/mic.20180344.

Tim J. Boonen and Mario Ghossoub. Bilateral risk sharing with heterogeneous beliefs and exposure constraints. *ASTIN Bulletin*, 50(1):293–323, 2020. doi:10.1017/asb.2019.39.

Kristof Bosmans and Z. Emel Öztürk. An axiomatic approach to the measurement of envy. *Social Choice and Welfare*, 50(2):247–264, 2018. doi:10.1007/s00355-017-1082-4.

Walter Bossert and Hans Peters. Single-peaked choice. *Economic Theory*, 41(2):213–230, 2009. doi:10.1007/s00199-008-0393-1.

Peter Caradonna and Zachary Raines. Interpersonally comparable utility. Working paper, August 25; arXiv:2608.25043. doi:10.2139/ssrn.7357439, 2026.

Jimmy Charité, Raymond Fisman, Ilyana Kuziemko, and Kewei Zhang. Reference points and redistributive preferences: Experimental evidence. *Journal of Public Economics*, 216:104761, 2022. doi:10.1016/j.jpubeco.2022.104761.

Janet Currie and Firouz Gahvari. Transfers in cash and in-kind: Theory meets the data. *Journal of Economic Literature*, 46(2):333–383, 2008. doi:10.1257/jel.46.2.333.

Terrence E. Daniel. A revised concept of distributional equity. *Journal of Economic Theory*, 11(1):94–109, 1975. doi:10.1016/0022-0531(75)90041-1.

Claude d'Aspremont and Louis Gevers. Equity and the informational basis of collective choice. *The Review of Economic Studies*, 44(2):199–209, 1977. doi:10.2307/2297061.

Claude d'Aspremont and Louis Gevers. Social welfare functionals and interpersonal comparability. In Kenneth J. Arrow, Amartya K. Sen, and Kotaro Suzumura, editors, *Handbook of Social Choice and Welfare*, volume 1, chapter 10, pages 459–541. Elsevier, 2002. doi:10.1016/S1574-0110(02)80014-5.

Bram De Rock and Domenico Moramarco. Multidimensional fairness with respect for individual preferences. *Social Choice and Welfare*, 66(4):883–909, 2026. doi:10.1007/s00355-025-01629-3.

Maya Eden. Welfare analysis with heterogeneous risk preferences. *Journal of Political Economy*, 128(12):4574–4613, 2020. doi:10.1086/710561.

Özgür Evren and Efe A. Ok. On the multi-utility representation of preference relations. *Journal of Mathematical Economics*, 47(4–5):554–563, 2011. doi:10.1016/j.jmateco.2011.07.003.

Marc Fleurbaey and François Maniquet. Fair allocation with unequal production skills: The no envy approach to compensation. *Mathematical Social Sciences*, 32(1):71–93, 1996. doi:10.1016/0165-4896(96)00811-6.
Marc Fleurbaey and François Maniquet. Fair income tax. *The Review of Economic Studies*, 73(1):55–83, 2006. doi:10.1111/j.1467-937X.2006.00369.x.
Marc Fleurbaey and François Maniquet. Fair social orderings. *Economic Theory*, 34(1):25–45, 2008. doi:10.1007/s00199-006-0132-4.
Marc Fleurbaey and François Maniquet. *A Theory of Fairness and Social Welfare*. Cambridge University Press, 2011. doi:10.1017/CBO9780511851971.
Marc Fleurbaey and François Maniquet. Fairness and well-being measurement. Preliminary and incomplete manuscript, March 2014, March 2014. URL https://pseweb.eu/ydepot/semin/MAN2015FAI.pdf.pdf.
Marc Fleurbaey and François Maniquet. Fairness and well-being measurement. *Mathematical Social Sciences*, 90: 119–126, 2017. doi:10.1016/j.mathsocsci.2016.11.003.
Marc Fleurbaey and François Maniquet. Inequality-averse well-being measurement. *International Journal of Economic Theory*, 14(1):35–50, 2018a. doi:10.1111/ijet.12140.
Marc Fleurbaey and François Maniquet. Optimal income taxation theory and principles of fairness. *Journal of Economic Literature*, 56(3):1029–1079, 2018b. doi:10.1257/jel.20171238.
Marc Fleurbaey and Koichi Tadenuma. Universal social orderings: An integrated theory of policy evaluation, inter-society comparisons, and interpersonal comparisons. *The Review of Economic Studies*, 81(3):1071–1101, 2014. doi:10.1093/restud/rdu006.
Marc Fleurbaey and Stéphane Zuber. Fair management of social risk. *Journal of Economic Theory*, 169:666–706, 2017. doi:10.1016/j.jet.2017.03.009.
Marc Fleurbaey and Stéphane Zuber. Fair utilitarianism. *American Economic Journal: Microeconomics*, 13(2):370–401, 2021. doi:10.1257/mic.20170234.
Duncan K. Foley. Resource allocation and the public sector. *Yale Economic Essays*, 7(1):45–98, 1967.
Leonardo Gasparini and Santiago M. Pinto. Equality of opportunity and optimal cash and in-kind policies. *Journal of Public Economics*, 90(1–2):143–169, 2006. doi:10.1016/j.jpubeco.2005.03.001.
Peter J. Hammond. Equity, arrow's conditions, and rawls' difference principle. *Econometrica*, 44(4):793–804, 1976. doi:10.2307/1913445.
Felix-Benedikt Liebrich. Risk sharing under heterogeneous beliefs without convexity. *Finance and Stochastics*, 28(4): 999–1033, 2024. doi:10.1007/s00780-024-00540-6.
R. Duncan Luce and John W. Tukey. Simultaneous conjoint measurement: A new type of fundamental measurement. *Journal of Mathematical Psychology*, 1(1):1–27, 1964. doi:10.1016/0022-2496(64)90015-X.
François Maniquet and Domenico Moramarco. Well-being measurement with reference consumption. *Economic Theory*, 82(2):437–476, 2026. doi:10.1007/s00199-026-01706-3.
François Maniquet and Yves Sprumont. Fair production and allocation of an excludable nonrival good. *Econometrica*, 72(2):627–640, 2004. doi:10.1111/j.1468-0262.2004.00503.x.
Michael Morreau and John A. Weymark. Measurement scales and welfarist social choice. *Journal of Mathematical Psychology*, 75:127–136, 2016. doi:10.1016/j.jmp.2016.04.001.
Hervé Moulin. One-dimensional mechanism design. *Theoretical Economics*, 12(2):587–619, 2017. doi:10.3982/TE2307.
Hervé Moulin and William Thomson. Can everyone benefit from growth? two difficulties. *Journal of Mathematical Economics*, 17(4):339–345, 1988. doi:10.1016/0304-4068(88)90016-X.
Elisha Pazner and David Schmeidler. Egalitarian equivalent allocations: A new concept of economic equity. *Quarterly Journal of Economics*, 92(4):671–687, 1978. doi:10.2307/1883182.
Paolo Giovanni Piacquadio. A fairness justification of utilitarianism. *Econometrica*, 85(4):1261–1276, 2017. doi:10.3982/ECTA14151.
John E. Roemer. *Theories of Distributive Justice*. Harvard University Press, 1996.
Emmanuel Saez and Stefanie Stantcheva. Generalized social marginal welfare weights for optimal tax theory. *American Economic Review*, 106(1):24–45, 2016. doi:10.1257/aer.20141362.
Norihito Sakamoto. No-envy, efficiency, and collective rationality. *Social Choice and Welfare*, 40(4):1033–1045, 2013. doi:10.1007/s00355-012-0657-3.
Karl H. Schlag. Identifying fairness with no envy: An axiomatization. SSRN working paper, posted August 2026. doi:10.2139/ssrn.7174659, 2026.
Amartya K. Sen. Choice functions and revealed preference. *The Review of Economic Studies*, 38(3):307–317, 1971. doi:10.2307/2296384.
Amartya K. Sen. Personal utilities and public judgements: Or what's wrong with welfare economics? *The Economic Journal*, 89(355):537–558, 1979. doi:10.2307/2231867.
Itai Sher. Generalized social marginal welfare weights imply inconsistent comparisons of tax policies. *American Economic Review*, 114(11):3551–3577, 2024. doi:10.1257/aer.20211025.
Yves Sprumont. Relative utilitarianism under uncertainty. *Social Choice and Welfare*, 53(4):621–639, 2019. doi:10.1007/s00355-019-01202-9.
William Thomson. Fair allocation rules. In Kenneth J. Arrow, Amartya K. Sen, and Kotaro Suzumura, editors, *Handbook of Social Choice and Welfare*, volume 2, chapter 21, pages 393–506. Elsevier, 2011. doi:10.1016/S0169-7218(10)00021-3.
U.S. Department of Agriculture, Food and Nutrition Service. The effects of cash-out on food use of food stamp

participants: Results from four demonstrations. Technical report, Office of Research and Analysis, September 1993. URL https://www.fna.usda.gov/research/snap/effects-cash-out-food-use-food-stamp-participants-results-four-demonstrations.

U.S. Department of Agriculture, Food and Nutrition Service. Special supplemental nutrition program for women, infants, and children (WIC): Revisions in the WIC food packages. Final Rule, 89 Federal Register 28488, 2024. URL https://www.federalregister.gov/d/2024-07437. FNS–2022–0007; RIN 0584–AE82.

Hal R. Varian. Equity, envy, and efficiency. *Journal of Economic Theory*, 9(1):63–91, 1974. doi:10.1016/0022-0531(74)90075-1.

Daniel Waldinger. Targeting in-kind transfers through market design: A revealed preference analysis of public housing allocation. *American Economic Review*, 111(8):2660–2696, 2021. doi:10.1257/aer.20190516.

Robert Wilson. The theory of syndicates. *Econometrica*, 36(1):119–132, 1968. doi:10.2307/1909607.

A. PROOFS OF THE MAIN RESULTS

These proofs follow the questions the main text answers. First we construct two nearby problems in which each has a different uniquely fair allocation but both offer the *same two own-welfare reports*. We then use the ordered recipient paths to prove the two cutoff inequalities and show how cost changes test the CCE references. Next we derive the peak rule, fit it with a CCE mixture, and characterize when maximin reliably selects the crossing. Finally we track the actual packages, budgets, cross-valuations, and transfer bounds in the reform and risk-sharing results. Each proof distinguishes evaluations of *one person's own package* from evaluations of *the other person's package*.

A.1. Fixed utility vectors constructed from own assignments

The cash–quality table gives one numerical reversal. The general impossibility theorem needs the same conflict for *every fixed preference-representing pair*. We first choose two nearby common packages x and y, with person 1 preferring y and person 2 preferring x. Around each we construct a line segment with a unique envy-free equal-package center. We then choose the opposite endpoint on each segment to have the *other center's own-welfare report*. Strict preference and indifference, not the utility scale, make the reports swap. A separate resource-balanced refinement, proved in the Additional Material, shows that smoothness and independent gradients can make both segments conserve aggregate resources.

Proof of Theorem 1. Fix a sufficiently small neighborhood of (x, x). We want two equal-package fair centers and two other allocations that swap the centers' reported welfare vectors. Choose a ball B around x such that $B \times B$ lies in the neighborhood and both preferences are continuous and strictly monotone there. Continuous utility representations U_1, U_2 exist on a suitable compact sub-ball. Local disagreement supplies a nearby y that person 1 strictly prefers to x but person 2 strictly dislikes relative to x, after interchanging the persons if needed.

Take a direction $r \gg 0$, which raises every component of an assignment. Choose $\bar{\sigma} > 0$ small enough to keep the assignments $x \pm 2\bar{\sigma}r$ inside B. Strict monotonicity gives strict utility inequalities on either side of x:

$$U_i(x - \bar{\sigma}r) < U_i(x) < U_i(x + \bar{\sigma}r), \qquad i = 1, 2.$$

By continuity, all assignments y sufficiently close to x also satisfy

$$U_1(y) < U_1(x + \bar{\sigma}r), \qquad U_2(x - \bar{\sigma}r) < U_2(y),$$

and

$$U_1(y - \bar{\sigma}r) < U_1(x), \qquad U_2(x) < U_2(y + \bar{\sigma}r).$$

The hypothesis supplies such a y with $y \succ_1 x$ and $x \succ_2 y$; choose it close enough that $y \pm \bar{\sigma}r \in B$.

We now use the inequalities to find the four *indifference corrections*. Along a positive direction r, each person's utility changes continuously and strictly upward. The intermediate value theorem therefore gives unique small movements $\sigma, \sigma', \tau, \tau' \in (0, \bar{\sigma})$ that make the following four comparisons bind:

$$x + \sigma r \sim_1 y, \qquad x - \sigma' r \sim_2 y, \qquad y - \tau r \sim_1 x, \qquad y + \tau' r \sim_2 x.$$

Define

$$G = \{(x + t\sigma r, x - t\sigma' r) : t \in [-1, 1]\},$$

and

$$H = \{(y - t\tau r, y + t\tau' r) : t \in [-1, 1]\}.$$

As the parameter moves along either segment, one recipient's package improves in every component while the other's worsens. Hence every allocation on each segment is Pareto efficient, and the paths satisfy the ordered-curve definition after the stated reparameterization. In G, person 1 envies to the left of $t = 0$ and person 2 to the right, so (x, x) is uniquely envy free. In H the roles of the two sides are reversed, and (y, y) is uniquely envy free. What remains is to check the two reported vectors at the opposite endpoints.

Fix any preference-representing indices e_1, e_2. Let ξ be the reported own-utility vector at the equal-package allocation (x, x), and let ξ' be the one at (y, y):

$$\xi = (e_1(x), e_2(x)), \qquad \xi' = (e_1(y), e_2(y)).$$

They are distinct because the two people rank x and y strictly in opposite directions. At $t = 1$ in G, the four indifferences above imply

$$W_e(x + \sigma r, x - \sigma' r) = \xi',$$

while at $t = 1$ in H,

$$W_e(y - \tau r, y + \tau' r) = \xi.$$

The four indifferences put ξ and ξ' in *both* own-report possibility sets. Yet efficient envy-free selection requires ξ in G and ξ' in H, because each segment has only its common-package center as an envy-free point. Each center's reported profile is unique along its segment by strict monotonicity. The same two reports would therefore be selected in opposite orders while both remain available, contradicting WARP. Since only indifference equalities were used to match the reports, changing either person's numerical utility scale cannot resolve the conflict. □

A.2. Interpersonal welfare comparisons implied by no-envy

The next proofs establish what the reader can infer from the ordered curve. As person 1 receives successively better packages and person 2 worse packages, person 1 eventually stops envying and person 2 eventually starts. We prove the resulting two cutoffs, then translate the inequalities into recipient-based CCE levels. The single-pair exception follows from the fact that one preference-representing pair can cross only once. Finally, the relative-cost proof shows how different costs make every package on an unchanged recipient path take its turn in a binding envy comparison.

Proof of the fair-set formula. Along the ordered curve, person 1's value of her own package rises while her value of person 2's falls, so her own-minus-other utility difference rises strictly with t. Person 2's corresponding difference falls strictly. The endpoint comparisons guarantee that each difference changes sign or reaches zero. Continuity then gives a unique cutoff for each, and the two signs yield $\gamma_G([a_G, b_G])$ when $a_G \leq b_G$, with no efficient envy-free allocation when $a_G > b_G$. □

Proof of Theorem 2. For part (a), each evaluator ranks the other recipient's path strictly, and an assignment's CCE is the position of its indifferent match on that path. Comparing this position with the recipient's actual position gives exactly (1) and (2); the identity $\pi_i^k(C_k(r)) = r$ handles each recipient's own assignment.

For part (b), agreement on C_k defines a common continuous strictly increasing function $h(r) = e_i(C_k(r))$. By Range, each $x \in D$ is indifferent for evaluator i to $C_k(\pi_i^k(x))$. Preference representation therefore gives

$$e_i(x) = e_i(C_k(\pi_i^k(x))) = h(\pi_i^k(x)).$$

Conversely, this factorization implies agreement on the path because $\pi_i^k(C_k(r)) = r$ for both evaluators.

For part (c), one pair of preference-representing indices must move in opposite directions as the curve shifts resources from person 2 to person 1. The difference between the two reports can therefore cross zero at most once. We ask whether this single crossing can reproduce *two distinct* envy equalities. Write the report difference as

$$\Delta_w(t) = w_1(C_1(t)) - w_2(C_2(1 - t)).$$

The first term is strictly increasing and the second strictly decreasing, so Δ_w is strictly increasing. Suppose the pair represents both directed comparisons. To represent person 1's test, the strictly increasing difference must be nonnegative exactly for $t \geq a_H$; to represent person 2's test, it must be nonpositive exactly for $t \leq b_H$. If $a_H < b_H$, both weak inequalities would hold at two different coordinates and force the strictly increasing difference to vanish twice. If $a_H > b_H$, any coordinate between the boundaries would require the difference to be strictly negative and strictly positive simultaneously. Hence both tests can be represented by one pair only if $a_H = b_H$. This argument uses strict monotonicity, not continuity of w.

Conversely, if $a_H = b_H = c$, the pair (π_1^2, π_2^2) has a strictly increasing reported-level difference with its unique zero at c. Its two opposing inequalities therefore reproduce the tests $t \geq c$ and $t \leq c$. For distinct cutoffs, part (a) supplies two pairs, while the preceding argument rules out a single pair. □

Proof of Proposition 4. To make a chosen person-1 package $T_1(s)$ appear at a binding envy equality, keep person 1's production path fixed and vary only person 2's relative production cost. Because both recipients strictly prefer more of their input, any efficient allocation must use the full resource budget. If person 1 uses input s and person 2 input t, the constraint is:

$$s + \kappa t = B.$$

Hence recipient 1's efficient assignments are exactly $T_1([0, B])$, independently of κ. Fix $s \in (0, B)$. The common endpoint classes and strict ordering give a unique $q_2(s) \in (0, B)$ such that

$$T_2(q_2(s)) \sim_2 T_1(s).$$

With

$$\kappa(s) = \frac{B - s}{q_2(s)} > 0,$$

the allocation $(T_1(s), T_2(q_2(s)))$ lies on the efficient frontier of $K_{1,\kappa(s)}$ and is exactly person 2's binding no-envy comparison. Since s was arbitrary, the boundary traverses the full interior of $T_1([0, B])$. The argument for $K_{2,\kappa}$ is symmetric: for each $t \in (0, B)$, let $q_1(t)$ satisfy $T_1(q_1(t)) \sim_1 T_2(t)$ and choose

$$\kappa(t) = \frac{B - t}{q_1(t)}.$$

This places person 1's binding comparison at $T_2(t)$. □

Proof of Theorem 3. Consider the family that keeps recipient 1's path unchanged. For each interior $T_1(s)$, the preceding proposition picks a cost at which person 2's own assignment is equally good to her as $T_1(s)$. If *the same indices* w_1, w_2 represent every such comparison, they must report a common level for this unchanged recipient package at each chosen cost. The following equalities make that obligation explicit.

$$T_2(q_2(s)) \sim_2 T_1(s)$$

and this indifference is person 2's no-envy boundary on the efficient frontier. If the fixed pair w represents that directed comparison, the binding equality and person 2's indifference give

$$w_1(T_1(s)) = w_2(T_2(q_2(s))) = w_2(T_1(s)).$$

The equality holds for every $s \in (0, B)$ and, by continuity, at both endpoints.

Every interior package on the fixed path has now appeared in a binding comparison, and continuity covers the two endpoints. Normalize the path as $C_1(r) = T_1(Br)$, and call its *common reported level*

$$h_1(r) = w_1(C_1(r)).$$

Both people strictly rank C_1 upward, so h_1 is continuous and strictly increasing and the path equality gives $w_i(C_1(r)) = h_1(r)$ for both i. Under Range, every $x \in D$ satisfies

$$x \sim_i C_1(\pi_i^1(x)).$$

Therefore

$$w_i(x) = w_i(C_1(\pi_i^1(x))) = h_1(\pi_i^1(x)),$$

which proves necessity. Conversely, suppose $w_i = h_1 \circ \pi_i^1$ on D, with h_1 continuous and strictly increasing. For each tested s,

$$w_1(T_1(s)) = h_1(s/B) = w_2(T_2(q_2(s))),$$

since $T_2(q_2(s)) \sim_2 T_1(s)$. Moreover, for every $x \in D$,

$$w_2(x) \geq w_1(T_1(s)) \quad \Longleftrightarrow \quad \pi_2^1(x) \geq s/B \quad \Longleftrightarrow \quad x \succsim_2 T_1(s).$$

Strict monotonicity of h_1 preserves the order of equivalent positions, so the last equivalence reproduces person 2's actual no-envy comparison. Repeating the argument with the second cost family proves the result for person 1's directed comparison. The domain remains D, where indifferent matches are known to exist; the proof does not claim that every package feasible at every test cost satisfies Range. □

A.3. Single-peaked choice and welfare aggregation

The choice proofs answer three separate questions. First, do consistent selections from *all intervals of one curve* reveal a peak? Second, can fixed calibrated CCE pairs be mixed so that their equality occurs at a selected fair peak? Third, which aggregator keeps that equality as the unique choice when the opposing indices change slope without moving their common level? The first question uses revealed choices, the second uses the two envy representations, and the third characterizes maximin. None alone supplies a representation of choices from disconnected menus.

Proposition 9 (Contraction consistency on one contract curve yields single-peaked choice)**.** *Under the restriction-domain conditions in Section 5, a single-valued rule that gives the same choice to problems with the same Pareto-efficient set satisfies Sen's condition α (contraction consistency) and continuity of the chosen coordinate on both one-sided interval families if and only if there is a unique point $\tau_H \in [0,1]$ such that every restriction G of H described in Section 5 satisfies (5). Thus τ_H is the peak of an ordinary single-peaked ordering on the contract curve.*

Proof of Proposition 9. Fix the parent problem H. Consider first the initial intervals $[0, r]$. Write $c_-(r)$ for the chosen coordinate and $\tau_H = c_-(1)$ for the full-curve choice. Sen's condition α says that whenever an interval still contains τ_H, it must still select τ_H. Thus

$$c_-(r) = \tau_H \qquad \text{for every } r \geq \tau_H,$$

because the point chosen from the full frontier remains feasible on each such initial restriction.

Suppose instead that an initial interval ending before the peak, $[0, r_0]$ with $r_0 < \tau_H$, chooses an interior $s < r_0$. Contraction consistency makes every smaller interval $[0, r]$ with $s \leq r \leq r_0$ choose the same s. To reach the full-curve choice τ_H as the right endpoint expands, the selected coordinate would eventually have to depart from s. One-sided continuity and contraction consistency together prevent that departure while s remains feasible.

If $c_-(r) = s < r < \tau_H$, continuity keeps the choice below r for r' just above r. The old choice is then still feasible on the larger interval, so contraction consistency forces $c_-(r') = s$. Thus the set of endpoints between r_0 and τ_H where choice equals s is nonempty, relatively open, and relatively closed. Since an interval is connected, this set extends toward τ_H; continuity gives $c_-(\tau_H) = s \neq \tau_H$, a contradiction. Therefore

$$c_-(r) = \min\{r, \tau_H\}.$$

The symmetric argument for terminal restrictions $[r, 1]$ gives

$$c_+(r) = \max\{r, \tau_H\}.$$

For an arbitrary retained interval $J = [\ell, u]$, condition α and the requirement that choice depend only on the Pareto-efficient set therefore imply that choice is the projection of τ_H onto J. Uniqueness follows from

choice on the baseline problem itself. On a compact interval this projection is exactly the unique maximizer of $-|r-\tau_H|$, giving the single-peaked representation in (5).

Conversely, if choice on every restriction of the same contract curve maximizes $-|r-\tau_H|$, menus with the same Pareto-efficient set have the same choice, and a chosen point remains chosen whenever a smaller feasible set retains it. Moreover, $\text{proj}_{[\ell,u]}\,\tau_H = \max\{\ell, \min\{\tau_H, u\}\}$ is jointly continuous in both endpoints. Thus projection implies the two one-sided continuity requirements and, more strongly, continuity on all retained intervals. □

The same projection rule also covers peaks at the two endpoints. With $\tau_H = 0$, every initial interval includes the peak and a terminal interval selects its left endpoint; with $\tau_H = 1$, reverse the roles. For an interior peak, the open-and-closed argument prevents choices from becoming stuck at an earlier interior point. The domain includes every pure-curve interval $\gamma_H(J)$, so the intervals used in that argument really are available. Pareto-set dependence then extends the result to restrictions containing additional dominated allocations. No numerical social welfare function is needed to establish projection.

In particular, the necessity proof uses every initial and terminal interval and continuity of choice on those one-sided families, but it does not require continuity on all two-sided intervals; projection itself supplies that stronger property. Without continuity or the full interval domain, contraction consistency alone would not force the result.

Proof of Proposition 5. Fix the two calibrations θ_1, θ_2 *before* using the observed choice. Each recipient's CCE on either path represents her preferences, so any positive mixture does as well. Along the ordered curve, person 1's mixed report rises with t and person 2's falls. What changes with α is the point at which these two reported levels meet. The proof locates that point by comparing the directed gaps.

$$d_k^H(t) = \theta_k\left(\pi_1^{k,H}(C_1^H(t))\right) - \theta_k\left(\pi_2^{k,H}(C_2^H(1-t))\right), \qquad k = 1, 2.$$

Each d_k^H is continuous and strictly increasing. Theorem 2 gives its unique zero at b_H for $k = 1$ and a_H for $k = 2$. The mixed difference is

$$d_\alpha^H(t) = \alpha d_1^H(t) + (1-\alpha) d_2^H(t).$$

It has exactly one zero $t_\alpha^H \in [a_H, b_H]$, since it is strictly increasing, nonpositive at a_H, and nonnegative at b_H.

Inside a nondegenerate fair interval, the difference from the recipient-1-based pair is negative and that from the recipient-2-based pair positive. To make the mixture difference zero at any chosen t in this interval, solve the single linear equation in α:

$$\alpha_H(t) = \frac{d_2^H(t)}{d_2^H(t) - d_1^H(t)}, \qquad t \in [a_H, b_H]. \tag{7}$$

The denominator is strictly positive throughout the closed fair interval. The formula includes $\alpha_H(a_H) = 0$ and $\alpha_H(b_H) = 1$. If $\alpha' < \alpha$, the mixed difference at $t_{\alpha'}^H$ is negative for weight α, so $t_{\alpha'}^H < t_\alpha^H$; the same argument covers the endpoints. For any convergent sequence of weights, every convergent subsequence of their zeros solves the limiting equality. Its uniqueness establishes continuity. Thus the zero map is a continuous strictly increasing bijection onto $[a_H, b_H]$, with inverse (7). If $a_H = b_H$, both directed differences vanish at the same point and every mixture has that unique zero. □

An index equality at the chosen coordinate fixes only one numerical comparison. Other strictly increasing transformations may preserve each person's preferences and the equality while changing all the reported gaps elsewhere. We first identify these equality-preserving transformations, then ask which social aggregator selects the crossing regardless of their opposing slopes. This separates invariance of the *one common level* from invariance under arbitrary renormalization.

Proposition 10 (Ordinal representations of the same welfare equality)**.** *Fix a continuous pair $z = (z_1, z_2)$ along the ordered contract curve such that z_1 represents person 1's ranking of her own assignments and is strictly increasing in t, z_2 represents person 2's ranking of her own assignments and is strictly decreasing in t, and*

$$z_1(\gamma_{H,1}(\tau)) = z_2(\gamma_{H,2}(\tau)) = c$$

only at τ. A second continuous pair $\widetilde{z} = (\widetilde{z}_1, \widetilde{z}_2)$ represents the same two individual rankings and has the same unique equality allocation τ if and only if there are continuous strictly increasing functions h_i, defined on the ranges of z_i, such that

$$\widetilde{z}_i(\gamma_{H,i}(t)) = h_i\big(z_i(\gamma_{H,i}(t))\big), \qquad i = 1, 2,$$

and $h_1(c) = h_2(c)$. Every such pair also agrees on whose measured welfare is lower on each side of τ.

Proof of Proposition 10. Each $z_i \circ \gamma_{H,i}$ is a continuous, strictly monotone function of t, so it is a homeomorphism onto its range. Any other continuous representation of the same ranking consequently differs by a continuous strictly increasing transformation

$$h_i = (\widetilde{z}_i \circ \gamma_{H,i}) \circ (z_i \circ \gamma_{H,i})^{-1}.$$

At the unique common level c, equality of the new pair is exactly $h_1(c) = h_2(c)$. Conversely, this condition gives equality at τ. For $t < \tau$, the original reports satisfy $z_1(\gamma_{H,1}(t)) < c < z_2(\gamma_{H,2}(t))$; strict monotonicity and the common transformed value preserve this ordering. The inequalities reverse for $t > \tau$, so the new pair has the same unique equality and identifies the same lower report on both sides. □

Proof of Proposition 6. The universal equality-selection property implies the same property for opposing linear indices under independent positive changes of slope. We show that even this restricted class forces the aggregator to depend only on the *lower* reported level. To see necessity, choose $a < b$ and an index rising just below a, while a second index falls to meet it at a. Specifically, let

$$x_\varepsilon(t) = a - \varepsilon(1 - t), \qquad y(t) = b - (b - a)t.$$

Take a compact interval containing 0 and having $t = 1$ in its interior. The first coordinate rises strictly, the second falls strictly, and their unique equality is at $t = 1$. The restricted equality-selection requirement therefore gives

$$W(a, a) > W(a - \varepsilon, b).$$

Letting $\varepsilon \downarrow 0$ and using continuity gives $W(a, a) \geq W(a, b)$, while coordinatewise monotonicity gives the reverse inequality. Thus raising a coordinate that is already above the minimum cannot increase W. Swapping the roles of the coordinates yields

$$W(z_1, z_2) = W(m, m), \qquad m = \min\{z_1, z_2\}.$$

Defining $\phi(m) = W(m, m)$ yields

$$W(z_1, z_2) = \phi(\min\{z_1, z_2\}),$$

and continuity and strict diagonal monotonicity make ϕ continuous and strictly increasing.

For the converse, any aggregator that is a strictly increasing transformation of the minimum selects the unique crossing. Before the crossing the smaller, increasing index determines the minimum; afterward the smaller, decreasing index does. Hence equality is the unique maximum, independently of the slopes. This completes the equivalence. □

Proof of Theorem 4. Suppose the behavioral conditions hold. Proposition 9 gives the projection rule (5); applying it to the full curve and imposing envy-free selection puts the revealed peak τ_H inside $[a_H, b_H]$. For the previously fixed calibrations, Proposition 5 then provides a mixture weight whose two component scores are equal at τ_H. To the left of that crossing person 1's mixed index is the lower score and rises; to its right person 2's mixed index is lower and falls. Their minimum is therefore maximized at the crossing and, on a restricted interval, at its nearest available coordinate. Thus one maximin score reproduces *every* interval choice required by the rule.

Conversely, every mixture has its equality point $t^H_{\alpha_H}$ in $[a_H, b_H]$, and its lower envelope is strictly single-peaked there, so maximin on an interval gives projection onto that point. Proposition 9 supplies single-valuedness, Pareto-set dependence, contraction consistency, and one-sided continuity. If $J_G \cap [a_H, b_H] \neq \emptyset$, the projection of a point of $[a_H, b_H]$ onto J_G belongs to that intersection, so the selected allocation is Pareto

efficient and envy free whenever such an allocation is available, proving *(i)*. Adding dominated allocations cannot change this argument because maximization is explicitly over $\mathrm{PE}(G)$.

The full-curve choice fixes τ_H. For a nondegenerate fair interval, the mixture weight is then unique *conditional on the calibrations*, by (7); at coincident boundaries, every weight gives the same choice. The universal aggregator proposition is separate: a particular rule on one problem does not identify maximin without its additional invariance requirement. □

A.4. Structural change and public benefits

The final proofs keep the actual *packages* in view. For inherited measurement we ask whether each new package has one old counterpart that both evaluators regard as indifferent. For one-sided improvement we track how each recipient's comparison changes at the original cutoff. For compensated benefits we check, in order, own utility, public spending, feasibility, and the two changed cross-valuations. The CARA transfer comparison is proved separately in the final subsection.

Portability criterion. To preserve a common welfare scale, each new recipient assignment must be indifferent, for both people, to the same old reference assignment. Suppose the baseline pair agrees along the old reference path, with common increasing level function h_0:

$$e_1(C_k^0(r)) = e_2(C_k^0(r)) = h_0(r).$$

If $C_k^1(r) \sim_i C_k^0(\rho(r))$ for both people, preference representation gives $e_i(C_k^1(r)) = h_0(\rho(r))$ for $i = 1, 2$, so the indices also agree on the new path.

For the converse, suppose the same indices also agree on the new path, with common level function h_1. If the *complete ranges* of h_0 and h_1 coincide, the strictly increasing inverse h_0^{-1} exists on the new range. Set $\rho = h_0^{-1} \circ h_1$, and for every r, the new and old packages then receive the same value from each preference-representing index:

$$e_i(C_k^1(r)) = h_1(r) = h_0(\rho(r)) = e_i(C_k^0(\rho(r))),$$

so $C_k^1(r) \sim_i C_k^0(\rho(r))$ for both people. Continuity and strict monotonicity make ρ a common increasing transformation of the path parameter. □

Proof of Proposition 7. At the old boundary a_0, person 1 was indifferent between her own assignment and person 2's. Improving person 2's package leaves person 1's assignment unchanged but makes the compared package weakly better:

$$C_1(a_0) \sim_1 C_2^0(1 - a_0).$$

The path improvement gives

$$C_2^1(1 - a_0) \succsim_1 C_2^0(1 - a_0),$$

so person 1 weakly prefers recipient 2's new assignment to her own assignment at $t = a_0$. On an ordered contract curve person 1 does not envy exactly for $t \geq a_1$. Hence $a_1 \geq a_0$. If the displayed path improvement is strict at $1 - a_0$, person 1 strictly envies at $t = a_0$, and uniqueness of the crossing gives $a_1 > a_0$.

The other boundary moves for a complementary reason. At b_0, person 2 was indifferent between her own and person 1's assignments; improving her own makes her weakly less inclined to envy:

$$C_2^0(1 - b_0) \sim_2 C_1(b_0),$$

and the path improvement gives

$$C_2^1(1 - b_0) \succsim_2 C_2^0(1 - b_0).$$

Person 2 therefore weakly prefers her own new assignment to person 1's at $t = b_0$. Since person 2 does not envy exactly for $t \leq b_1$, it follows that $b_1 \geq b_0$, with strict inequality when the path improvement is strict for person 2 at $1 - b_0$.

Finally, if $\tau_0 < a_1$, the fair-set characterization implies that person 1 strictly envies at $\gamma^1(\tau_0)$. The first component of that allocation is $C_1(\tau_0)$, exactly as before the reform. If the new fair set is nonempty, every efficient envy-free allocation has parameter in $[a_1, b_1]$, and hence at least $a_1 > \tau_0$. □

Singleton policy reversal. Because the old fair set is the singleton $\{t_0\}$, $a_0 = b_0 = t_0$. The strict improvement at $1 - t_0$ gives $a_1 > a_0$ by Proposition 7. Since the new fair set is the singleton $\{t_1\}$, $t_1 = a_1 > t_0$. On the old ordered contract curve, person 2 does not envy exactly for $t \leq b_0 = t_0$, so person 2 strictly envies at t_1. On the new curve, person 1 does not envy exactly for $t \geq a_1 = t_1$, so person 1 strictly envies at t_0. The stated choice reversal follows. □

Access formulas and Pareto comparison. At $\gamma^\kappa(t)$, person i's utility from person 1's assignment is $(1 + \beta_i)t$, whereas utility from person 2's assignment, with $r = 1 - t$, is

$$(1 + \beta_i\kappa)r + \beta_i(1 - \kappa)r^2.$$

Put

$$c_i^\kappa = (1 - \kappa)\frac{\beta_i}{1 + \beta_i}.$$

The binding comparison is equivalent to

$$t = (1 - t)(1 - c_i^\kappa t),$$

or

$$2t - 1 + c_i^\kappa t(1 - t) = 0.$$

Its unique root in $[0, 1]$ is

$$f(c_i^\kappa) = \frac{2}{2 + c_i^\kappa + \sqrt{4 + (c_i^\kappa)^2}}.$$

For person 1 this is the lower boundary a_κ; for person 2 it is the upper boundary b_κ. The function f is strictly decreasing on $[0, \infty)$, while c_i^κ strictly decreases in $\kappa < 1$. Hence both boundaries rise. Since

$$\frac{\beta_1}{1 + \beta_1} > \frac{\beta_2}{1 + \beta_2},$$

we have $a_\kappa < b_\kappa$ for $\kappa < 1$, and both equal $1/2$ at $\kappa = 1$. The one-versus-two statement follows from Theorem 2.

Finally, $a_{\kappa_1} > b_{\kappa_0}$ if and only if

$$(1 - \kappa_1)\frac{\beta_1}{1 + \beta_1} < (1 - \kappa_0)\frac{\beta_2}{1 + \beta_2},$$

because f is strictly decreasing. Rearranging yields

$$\kappa_1 > 1 - (1 - \kappa_0)\frac{\beta_2(1 + \beta_1)}{\beta_1(1 + \beta_2)}.$$

For the Pareto comparison, fix $\kappa_0 < 1$ and put $a = a_{\kappa_0}$. Every old fair allocation has $t < 1/2$, so person 1 strictly prefers $C_1(1/2)$ to her old assignment. Person 2's own utility is highest over the old fair interval at its lower endpoint $t = a$. The person-1 boundary equation at a implies, after substitution,

$$U_2(C_2^1(1/2)) - U_2(C_2^{\kappa_0}(1 - a)) = \frac{(1 - 2a)\left[\beta_1\beta_2 - \beta_1 + 2\beta_2\right]}{2\beta_1}.$$

Since $a < 1/2$, this difference is strictly positive exactly when

$$\beta_1(1 - \beta_2) < 2\beta_2.$$

Under that condition person 2 strictly prefers the full-access fair assignment to her most preferred assignment in the old fair interval, and hence to every old fair allocation. □

Proof of Theorem 5. The benefit theorem requires three distinct verifications. We first show that every recipient is indifferent between her old and new package at the same resource coordinate. We next check the public budget for *all feasible pairs of coordinates*, establishing that the own-utility possibility set is unchanged. Only then do we compare each person with the other's *new* package and derive the two potentially different envy cutoffs. Put $P = 1 + p$ and write the positive, budget-balancing adjustment A as:

$$A = (\beta_1 - p)\delta_1 = (p - \beta_2)\delta_2 > 0.$$

Own-utility neutrality follows directly:

$$U_i(C_i^\delta(r)) = r - \beta_i\delta_i v(r) + \beta_i\big(r + \delta_i v(r)\big) = (1 + \beta_i)r.$$

The aggregate public cost at indices (s, t) is

$$K_\delta(s, t) = P(s + t) - Av(s) + Av(t).$$

Symmetry gives $v(1 - s) = v(s)$, so $K_\delta(s, 1 - s) = P$. For fixed s, the map

$$t \longmapsto Pt + Av(t) = c_p(C_2^\delta(t))$$

is strictly increasing by assumption. Hence $K_\delta(s, t) \leq P$ if and only if $t \leq 1 - s$. The feasible index set is therefore $s + t \leq 1$, exactly as before the reform. Common strict ordering makes $s + t = 1$ the efficient frontier, and own-utility neutrality then gives $\mathcal{U}^0 = \mathcal{U}^\delta$.

The new packages give recipients the same utility at each coordinate and use the same feasible resource indices, yet the two *cross-package* values change. On the efficient frontier, person 1 must compare her own replacement with person 2's and person 2 must make the reverse comparison. These inequalities, rather than the own-utility identities, determine the new envy boundaries.

$$(1 + \beta_1)(2s - 1) \geq \Delta\delta_2 v(s),$$

and person 2 does not envy if and only if

$$(1 + \beta_2)(2s - 1) \leq \Delta\delta_1 v(s).$$

No $s \leq 1/2$ satisfies the first inequality after a nonzero reform. For $s > 1/2$, the two conditions are

$$\frac{\Delta\delta_2}{1 + \beta_1} \leq \psi_v(s) \leq \frac{\Delta\delta_1}{1 + \beta_2}.$$

Because ψ_v is strictly increasing from 0 to $+\infty$, the two boundary points are the displayed a and b. Efficient no-envy exists exactly when

$$\frac{\delta_2}{1 + \beta_1} \leq \frac{\delta_1}{1 + \beta_2}.$$

Using $\delta_2/\delta_1 = (\beta_1 - p)/(p - \beta_2)$, this is equivalent to

$$p \geq \frac{\beta_1 + \beta_2 + 2\beta_1\beta_2}{2 + \beta_1 + \beta_2} = p^*.$$

If the first cutoff is below the second, the interval between them is efficiently envy free; if they meet, only one coordinate is; if their order reverses, none is. The original economy has both cutoffs at $1/2$. The one-pair exception then follows from Theorem 2 *on the efficient curve*, not on all feasible package pairs. □

Robustness to imperfect compensation. Let $\bar{v} = \max_{r\in[0,1]} v(r)$ and define the perturbed packages, for sufficiently small $|\lambda|$, by

$$C_{1,\lambda}(s) = \big(s - \beta_1\delta_1 v(s) + \lambda v(s),\ s + \delta_1 v(s)\big), \qquad C_{2,\lambda}(t) = \big(t - \beta_2\delta_2 v(t) - \lambda v(t),\ t + \delta_2 v(t)\big).$$

Put $P = 1 + p$ and $A = (\beta_1 - p)\delta_1 = (p - \beta_2)\delta_2$. The two public costs are

$$Ps - (A - \lambda)v(s) \qquad \text{and} \qquad Pt + (A - \lambda)v(t).$$

Their nonlinear terms cancel at $t = 1 - s$. For sufficiently small $|\lambda|$, the individual cost schedules remain strictly increasing, so the public-budget inequality is equivalent to $s + t \leq 1$, and every efficient allocation exhausts the budget.

The intended-recipient utilities are

$$U_1(C_{1,\lambda}(s)) = (1 + \beta_1)s + \lambda v(s), \qquad U_2(C_{2,\lambda}(t)) = (1 + \beta_2)t - \lambda v(t).$$

Hence every perturbed utility vector lies within $|\lambda|\bar{v}$ in sup norm of its exactly compensated counterpart under the same indices, and conversely. This proves the Hausdorff bound.

On the efficient frontier, person 1's no-envy condition becomes

$$(1 + \beta_1)(2s - 1) \geq (\Delta\delta_2 - 2\lambda)v(s),$$

while person 2's becomes

$$(1 + \beta_2)(2s - 1) \leq (\Delta\delta_1 - 2\lambda)v(s).$$

For sufficiently small $|\lambda|$, both right-hand coefficients are positive. Strict monotonicity of ψ_v then gives the two equality boundaries

$$a_\lambda = \Psi_v\left(\frac{\Delta\delta_2 - 2\lambda}{1 + \beta_1}\right), \qquad b_\lambda = \Psi_v\left(\frac{\Delta\delta_1 - 2\lambda}{1 + \beta_2}\right).$$

At $\lambda = 0$, strict separation for $p > p^*$ and strict reversal for $p < p^*$ are open inequalities. Continuity therefore preserves the respective fair-set conclusions for all sufficiently small $|\lambda|$. □

Proof of Corollary 1. The corollary follows from *each person's own indifference*: at a fixed resource index she is indifferent between her original and replacement assignments,

$$C_i^\delta(r) \sim_i C_i^0(r).$$

Thus every fixed index e_i representing $\succsim_i$ on the two package paths satisfies

$$e_i(C_i^\delta(r)) = e_i(C_i^0(r)).$$

Together with the common feasible index set from Theorem 5, this gives

$$\mathcal{W}_e^0 = \mathcal{W}_e^\delta$$

for every fixed pair $e = (e_1, e_2)$; the equality of utility possibility sets is the special case $e_i = U_i$.

The reformed lower cutoff is strictly above the old equal share $1/2$ whenever the adjustment is nonzero. Thus the reformed fair set never includes the old fair share, and for $p < p^*$ it is empty. The *information* conclusion is strongest when $p > p^*$: both environments offer nonempty fair sets but require different fair choices. When $p < p^*$, no efficient fair allocation exists regardless of what the planner observes.

Suppose a single-valued rule depends only on $\mathcal{W}_e^E$ and implements the baseline's unique efficient envy-free allocation. The input set is unchanged after the reform, so the rule selects the same utility vector. Along each recipient's path, $e_i(C_i^E(r))$ is strictly increasing in r, and the compensation equalities match these values across environments, so the repeated profile corresponds to $s = t = 1/2$. If $p \geq p^*$, the reformed efficient envy-free set begins at $a > 1/2$, so person 1 envies at that retained choice, whereas if $p < p^*$, no efficient envy-free allocation exists. No such fixed own-assignment welfare rule implements efficient no-envy in both environments.

Finally, the baseline boundaries coincide at $1/2$, whereas $a < b$ after the reform whenever $p > p^*$. The one-pair-versus-two-pair conclusion follows directly from Theorem 2(c). □

Proof of Corollary 2. At a fixed index r, both the baseline and reformed package give intended recipient i utility $(1 + \beta_i)r$. Person 1's reformed package costs

$$Pr - Av(r),$$

less than the baseline package Pr, while person 2's baseline package costs less than person 2's reformed package. Giving person 1 the reformed package and person 2 the baseline package at a common index r therefore costs

$$2Pr - Av(r).$$

The budget equation $2Pr - Av(r) = P$ is equivalent to

$$\psi_v(r) = A/P,$$

so its unique solution is $r = \Psi_v(A/P) > 1/2$.

Person 1 is indifferent between the two assigned packages, because both give person 1 utility $(1 + \beta_1)r$. Person 2 strictly prefers her own baseline package: person 2 values person 1's reformed package at

$$(1 + \beta_2)r - \Delta\delta_1 v(r) < (1 + \beta_2)r.$$

The allocation is therefore envy free. It uses the least costly available package for each recipient at index r. These least-cost schedules are strictly increasing under the maintained smallness condition, so a Pareto improvement would require more than the available budget. The allocation is Pareto efficient. Since $r > 1/2$, both recipients strictly prefer it to the original fair allocation. □

A.5. CARA risk sharing and money-unit identification

The risk-sharing proof evaluates both contracts under each recipient's own certainty equivalent, obtains the two transfer bounds, and proves their strict order using convexity of the log moment-generating function. The final money-unit lemma identifies the additive corrections required when welfare reports must increase one for one with sure cash.

Proof of Theorem 6. The CARA proof first evaluates *both* risky shares from *each* recipient's perspective, then solves the two no-envy inequalities for the sure transfer d. The only non-obvious part is showing that the lower transfer bound is strictly below the upper. Efficient CARA sharing follows from standard first-order conditions; strict ordering of the bounds follows from the positive variance of nonconstant aggregate consumption under exponential reweighting of the common belief.

$$\mathrm{CE}_i(c_1) = A_i + d, \qquad \mathrm{CE}_i(c_2) = D_i - d, \qquad i = 1, 2.$$

Thus person 1 does not envy exactly when $d \geq (D_1 - A_1)/2 = d_1$, and person 2 does not envy exactly when $d \leq (D_2 - A_2)/2 = d_2$.

The remaining issue is whether the lower transfer bound d_1 is strictly below the upper bound d_2. Let $L(v)$ be the log moment-generating expression

$$L(v) = \log E_p[e^{-vB}], \qquad \gamma = 1 - \alpha.$$

Because B is nonconstant and p has full support, $L''(v)$ is the strictly positive variance of B under the exponentially tilted probability. Moreover,

$$D_\eta - A_\eta = \frac{L(\eta\alpha) - L(\eta\gamma)}{\eta} = \int_\gamma^\alpha L'(\eta z)\, dz.$$

Differentiating the integral gives $F'(\eta) = \int_\gamma^\alpha zL''(\eta z)\, dz$ for $F(\eta) = D_\eta - A_\eta$. Since B is nonconstant and beliefs have full support, $L'' > 0$. Hence F' has the sign of $\alpha - \gamma$, which is also the sign of $\eta_2 - \eta_1$. Moving from η_1 to η_2 therefore *strictly increases* F, whether $\eta_2 > \eta_1$ or $\eta_2 < \eta_1$:

$$D_2 - A_2 > D_1 - A_1.$$

so $d_2 > d_1$.

For the first directed pair,

$$z_1^2(c_1) = A_1 + d, \qquad z_2^2(c_2) = D_2 - d + D_1 - D_2 = D_1 - d,$$

so equality occurs at d_1. For the second,

$$z_1^1(c_1) = A_1 + d + A_2 - A_1 = A_2 + d, \qquad z_2^1(c_2) = D_2 - d,$$

so equality occurs at d_2. An affine mixture has a welfare difference that is strictly increasing in d and whose zero moves continuously from d_1 to d_2, hence spans the fair interval.

Maximin of the unadjusted own certainty equivalents equalizes A_1+d and D_2-d, giving $d^{\text{own}} = (D_2-A_1)/2$. If $\eta_1 < \eta_2$, person 1 is less risk averse, so $D_1 > D_2$ for the nonconstant lottery $(1-\alpha)B$, and $d^{\text{own}} < d_1$. If $\eta_1 > \eta_2$, person 2 is less risk averse, so $A_2 > A_1$, and $d^{\text{own}} > d_2$. The selected contract is envious in either case. □

Lemma 1 (Money-unit identification)**.** *Let a comparison domain be closed under sure cash translations and let each preference admit a cash-additive certainty equivalent* CE_i*. Fix a recipient path* $C_k(r) = b_k + r\mathbf{1}$*. If welfare indices represent the preferences, satisfy*

$$w_i(x + a\mathbf{1}) = w_i(x) + a,$$

and correctly represent one binding comparison with C_k *as welfare equality, then, for one common constant* K*,*

$$w_i(x) = \text{CE}_i(x) - \text{CE}_i(b_k) + K, \qquad i = 1, 2.$$

Conversely, this pair represents every directed comparison with the path.

Proof. A money-normalized index must increase by one whenever a sure dollar is added to any contract. Fix one assignment x_0 in the cash-translation domain. For another contract x, cash additivity makes x indifferent to $x_0 + (\text{CE}_i(x) - \text{CE}_i(x_0))\mathbf{1}$. Preference representation and the unit restriction therefore force $w_i(x)$ to equal its certainty equivalent plus a constant independent of x:

$$w_i(x) = w_i(x_0) + \text{CE}_i(x) - \text{CE}_i(x_0) = \text{CE}_i(x) + c_i,$$

where $c_i = w_i(x_0) - \text{CE}_i(x_0)$ does not depend on x. At a binding comparison $y \sim_j C_k(r_0)$, the required equality $w_j(y) = w_k(C_k(r_0))$ implies

$$\text{CE}_j(b_k) + c_j = \text{CE}_k(b_k) + c_k =: K$$

after subtracting the common cash amount r_0. Hence $c_i = K - \text{CE}_i(b_k)$, as claimed. Conversely, $\text{CE}_i(x) - \text{CE}_i(b_k)$ is exactly the cash coordinate of x on C_k, so the resulting pair represents every directed comparison with that path. □