Fair Policy Choice under Risk and Constraints: Certainty-Equivalent Complaints and a Veil Foundation*

Niels Boissonnet† and Yves Breitmoser‡

February 26, 2026

Abstract

Many distributive policy problems under risk are choices from constrained menus shaped by incentive, participation, balanced-budget, and institutional restrictions. These constraints can make envy-free efficiency unattainable even on the Pareto frontier, so a fairness criterion must specify how to choose among efficient but imperfect reforms. We develop an operational complaint-based rule for such environments. Each identity evaluates a policy by its baseline-relative certainty-equivalent (CE) net gain, a money-metric index we characterize by preference consistency, translation invariance, and calibration to sure shifts. The resulting CE-indexed selector chooses an efficient policy that equalizes CE net gains whenever feasible; otherwise it selects the Pareto-efficient policy that minimizes the maximum residual complaint in CE dollars, with a leximin refinement. We also provide a veil-of-ignorance foundation: the same choice minimizes worst-case ex post regret over identities in CE units. In a flagship application to adverse selection and insurance incidence, we characterize the fair contribution ceiling under a mandate, the fair pooled policy under voluntary enrollment when pooling is feasible, and a CE-dollar measure of the fairness loss from unraveling when pooling is infeasible.

Keywords: Fairness; envy; complaint-based choice; certainty equivalents; constrained policy menus; adverse selection; insurance incidence.

JEL codes: D63; D81; D82; H21; I13.

*This is a preliminary draft, we are grateful for all comments. Comments of Gerrit Bauch, Jurek Preker, Frank Riedel and the audience of the Economic Theory Lunch Seminar in Bielefeld are greatly appreciated. Financial support by Verbundinitiative InChange funded by MKW NRW is also greatly appreciated.

†Department of Economics, IMW, and JICE, Bielefeld University

‡Corresponding author. Department of Economics, IMW, and JICE, Bielefeld University

1 Introduction

Policy choice under risk is typically made from a constrained set of feasible instruments. A health-insurance reform must satisfy balanced-budget requirements and institutional restrictions on the contribution base; unemployment insurance must respect payroll-tax capacity and legal minima; credit policy is constrained by collateral and incentive-compatibility requirements. These constraints matter for distributive evaluation because they can make canonical fairness benchmarks infeasible even when the planner agrees on the benchmark itself. The normative problem is therefore not only to articulate a fairness ideal, but also to specify a selection rule that respects feasibility and efficiency and remains disciplined when the ideal cannot be attained.

A natural way to formalize fairness is in terms of complaints. In deterministic environments, envy-freeness has long served as a complaint-based benchmark. In economies with divisible goods, Varian (1974) studies envy-free allocations as a central fairness notion. In probabilistic assignment, Bogomolnaia and Moulin (2001) develop ordinal notions of fairness and highlight the tradeoffs created by feasibility constraints. In matching problems with priorities, Abdulkadiroğlu et al. (2020) propose justified envy as a feasibility-aware complaint notion. Bringing this complaint-based perspective to policy problems under risk requires a money-metric language: identities can differ both in risk exposure and in attitudes toward risk, so complaints must be expressed in a unit that remains comparable across individuals. Money-metric welfare analysis provides such a language (Blackorby and Donaldson, 1988), and certainty equivalents provide a natural projection of risky prospects into monetary units (Broome, 2012).

This paper studies fair policy choice from constrained menus under risk and uncertainty. A menu G is a set of feasible policies generated by incentive, participation, budget, and institutional constraints. Given a menu G and a baseline policy g^0 (the status quo), the planner must choose a policy $g \in G$ that is Pareto efficient and fair in the sense of limiting justified complaints. We build on the menu-based surrogate approach developed in our companion paper *Fair choice under risk and uncertainty*, which characterizes complaint-based selectors on constrained menus and provides a principled continuation along the Pareto frontier when envy-free efficiency is infeasible. The substantive step taken here is to make that approach operational in economics settings with risk by measuring complaints in certainty-equivalent money units, linking the resulting criterion to veil-of-ignorance reasoning, and developing policy applications where feasibility constraints are economically central.

Our main results can be stated without technicalities. We evaluate each identity's gain from a policy relative to the baseline in a common money-metric unit: the baseline-relative certainty-equivalent net gain. If the menu contains an efficient policy that equalizes these gains across identities, the fair choice rule selects an equalizer. If feasibility prevents equalization even on the Pareto frontier, the rule selects an efficient policy that makes the remaining complaints as small as possible in the same money-metric units, with a leximin refinement for tie-breaking. The same policy can also be characterized behind the veil of ignorance: it minimizes the worst-case ex post regret across identities when regret is evaluated in certainty-equivalent dollars. This links the continuation rule to canonical veil-of-ignorance

perspectives, including Rawls' impartiality argument (Rawls, 1971) and Harsanyi's aggregation argument (Harsanyi, 1955).

The next example illustrates why a continuation rule is needed in a workhorse setting and why the relevant unit is certainty-equivalent dollars.

Example 1 (Insurance incidence and the need for a continuation rule)**.** *Consider two groups $i \in \{L, H\}$ facing a binary loss $\ell > 0$ with probabilities $q_L > q_H$ and incomes $w_L < w_H$. A public insurance policy provides coverage $I \in [0, \ell]$ and charges deterministic contributions t_i. Relative to the baseline of no insurance, let $V_i(I)$ denote the certainty-equivalent value of coverage and define the certainty-equivalent net gain*

$$\Gamma_i \;=\; V_i(I) - t_i.$$

Absent moral hazard, full coverage $I = \ell$ is efficient, so the central distributive issue is incidence, namely who pays for pooled coverage. This incidence problem is the organizing concern of classic work on adverse selection (Akerlof, 1978) and of the modern empirical literature that measures demand, costs, and welfare consequences of selection in insurance markets (Einav and Finkelstein, 2011, 2023). A natural complaint-based benchmark in this setting is to equalize certainty-equivalent net gains across groups, $\Gamma_L = \Gamma_H$, so that the risk-adjusted gains from the reform are the same in money-metric terms.

The difficulty is that feasibility restricts incidence. Suppose contributions are proportional to an earnings base capped at a ceiling $\bar{w}$ (a contribution ceiling), with the rate pinned down by budget balance. As the ceiling moves, contributions shift between groups, generating a one-dimensional Pareto frontier of efficient mandated policies indexed by $\bar{w}$. Equalizing certainty-equivalent net gains requires an incidence condition of the form

$$t_H(\bar{w}) - t_L(\bar{w}) \;=\; V_L(\ell) - V_H(\ell).$$

If the feasible range of ceilings is narrow, or if legal constraints restrict the contribution base, this equality may have no solution. Even along the efficient frontier, certainty-equivalent net gains cannot be equalized. In that case, an envy-free benchmark is silent unless it is supplemented by a principled continuation rule that selects, among Pareto-efficient feasible policies, the one that makes the remaining complaints as small as possible. In the two-group case, this continuation is transparent: choose the frontier point that minimizes $|\Gamma_L - \Gamma_H|$, equivalently minimizing worst-case ex post regret behind the veil in certainty-equivalent dollars.

Our contributions are threefold. First, we provide a canonical complaint unit for policy choice under risk by studying baseline-relative certainty-equivalent net gains and showing how they can be computed and approximated using local objects familiar from applied welfare analysis. In particular, our approximation results connect to sufficient-statistics approaches to welfare evaluation (Finkelstein and Hendren, 2020) and to unified frameworks for policy evaluation (Hendren and Sprung-Keyser, 2020). Second, combining these money-metric gains with the companion paper's menu-based selector yields a sharp characterization of fair choice on constrained efficient sets: equalize certainty-equivalent net gains when

feasible, and otherwise minimize the maximal residual complaint in certainty-equivalent dollars with a leximin refinement. Third, we provide a veil-of-ignorance interpretation that selects the same policy by minimizing worst-case ex post regret across identities, which clarifies the content of the continuation rule and anchors it in standard normative thought experiments.

Section 3 formalizes the menu problem and takes the companion paper's surrogate selector as given. Sections 4–6 develop the certainty-equivalent gauge, the resulting fair choice rule, and its veil interpretation. Section 7 develops the insurance application, and Section 9 develops additional applications and extensions, including statutory tax reform.

2 Related literature and relation to the companion paper

This paper builds on a companion paper that develops a general framework for fair choice under risk and uncertainty from *menus* of feasible policies. The companion paper takes feasibility constraints as primitive, restricts attention to Pareto-efficient policies, and characterizes complaint-based selectors that coincide with envy-free efficiency whenever that benchmark is attainable and otherwise provide a disciplined continuation along the Pareto frontier. The present paper makes that framework directly usable for economics applications with risk. We fix a canonical interpersonal gauge using baseline-relative certainty-equivalent (CE) net gains, we show that the resulting selector admits a veil-of-ignorance interpretation in the same CE units, and we develop applications in which feasibility constraints are economically substantive. Section 3 states the menu problem and takes the surrogate selector as given, while Sections 4–6 supply the economic content that turns the selector into an operational criterion.

A first point of contact is the complaint-based tradition in distributive economics. Foley (1967) introduces envy-freeness as a fairness benchmark, and Varian (1974) gives it a competitive interpretation. Thomson (2011) surveys the broader fair-division literature that develops this complaint logic. Once feasibility constraints bind, complaint-free benchmarks can fail to exist even on the efficient set. The literature responds either by relaxing no-envy in ways that restore existence, as in the λ-equitability refinement of Diamantaras and Thomson (1990), or by ranking efficient allocations by the severity of envy, as in the "intensity of envy" ordering of Chaudhuri (1986). Related feasibility issues arise with indivisibilities, where approximate fairness preserves the complaint interpretation while guaranteeing existence; Budish (2011) provides an approximate competitive equilibrium from equal incomes for such environments. In matching settings with priorities, Abdulkadiroğlu et al. (2020) formalize "justified envy" as a feasibility-aware complaint notion. Our contribution fits this line of work but targets constrained *policy* menus under risk: the companion paper provides a general continuation principle for selecting among efficient feasible reforms when complaint-free benchmarks may be empty, and the present paper supplies economically meaningful units in which residual complaints are compared.

A second set of connections concerns money-metric and egalitarian-equivalent approaches to interpersonal comparability. d'Aspremont and Gevers (1977) and Pazner and Schmeidler (1978) separate the choice of a common metric from the choice of ethical aggregation, and Blackorby and Donaldson (1988) and Fleurbaey (2009) develop money-metric constructions for welfare comparisons with heterogeneous preferences. Under risk, certainty equivalents provide the canonical projection of uncertain prospects into monetary units, in the sense emphasized by Broome (2012). We build on this projection logic but anchor comparability in the policy problem through the baseline g^0: baseline-relative CE net gains measure incidence as "who gains and who pays" in a unit that respects each identity's own risk attitudes, which is the relevant comparability requirement for complaint-based evaluation in risky environments (Section 4).

A third strand is modern public finance and empirical welfare analysis. Optimal tax theory provides benchmark prescriptions for redistribution under incentive constraints; Mirrlees (1971) introduces the optimal income tax problem, Atkinson and Stiglitz (1976) analyze commodity taxation and separability conditions, and Saez (2001) develops implementable sufficient-statistics formulas for optimal taxes. More recent work organizes policy evaluation using generalized social marginal welfare weights; see Saez and Stantcheva (2016). In empirical work, sufficient-statistics approaches translate causal estimates into money-metric welfare criteria. Finkelstein and Hendren (2020) develop the marginal value of public funds framework, and Hendren and Sprung-Keyser (2020) propose a unified decomposition for policy evaluation. At the same time, weight-based local criteria can raise coherence concerns unless they are anchored in a globally consistent objective; Sher (2024) articulates this concern in a recent contribution. Our approach is complementary. We also work with money-metric objects, but we use baseline-relative CE net gains as the unit and the constrained efficient set as the discipline that organizes interpersonal comparisons in the complaint domain. We connect these CE units to the same empirical objects that underlie sufficient-statistics welfare analysis, while the selector pins down how to trade off identities' gains when feasibility limits what reforms can achieve.

The veil-of-ignorance interpretation in Section 6 links our analysis to debates about ex ante and ex post evaluation under uncertainty about identity. Harsanyi (1955) motivates ex ante aggregation of von Neumann–Morgenstern utilities via an impartial-observer argument, while Rawls (1971) emphasizes protecting the worst-off position behind the veil. A large welfare-theoretic literature studies tensions between efficiency, equity, and rationality under risk and uncertainty; Miyagishima (2022) provides recent contributions, and Fleurbaey and Zuber (2022) survey and extend prominent impossibility results. Our veil result differs in a way that matters for constrained policy evaluation. The object being compared is a CE net gain relative to a baseline policy, and the criterion is intrinsically feasibility-aware: when equalization is infeasible on the Pareto frontier, the rule selects the efficient feasible policy that minimizes worst-case ex post regret in CE dollars, with a leximin refinement. Related veil-of-ignorance interpretations for fairness criteria have also been developed in adjacent literatures; see Liang and Lu (2024a).

Finally, the insurance and tax applications connect our criterion to large literatures on selection, regulation, and incidence. The canonical adverse selection framework originates

in Akerlof (1970) and Rothschild and Stiglitz (1976). The modern empirical literature measures the welfare consequences of selection and studies the incidence effects of regulation and mandates; Einav and Finkelstein (2011) develop the central welfare tools, and Hackmann et al. (2015) analyze mandates and incidence in regulated markets. Our contribution is not to re-derive selection inefficiencies, but to provide a tractable fairness criterion for choosing among efficient feasible reforms when incidence is constrained by institutions and when groups differ in risk exposure and risk preferences (Section 7). Tax reform provides a parallel constrained-menu setting in which legal form, enforcement, and administrative limits restrict incidence even before behavioral responses are considered; our wealth-tax application illustrates how baseline-relative CE net gains can be used to evaluate progressive proposals within a statutory class, complementing quantitative work on efficiency and redistribution under wealth taxation such as Guvenen et al. (2023).

3 Menus, Pareto efficiency, and the surrogate selector

This section formalizes policy choice from constrained feasible sets and recalls the surrogate selector from our companion paper. The substantive economic content of the present paper enters only once we have fixed a meaningful unit in which complaints can be compared under risk. For now, we take feasibility constraints as given and focus on disciplined choice along the resulting Pareto frontier.

Complaint-based fairness benchmarks motivate the selection problem. Envy-freeness provides a canonical way to express complaints in deterministic environments (Foley, 1967), and Varian (1974) gives a competitive interpretation. In constrained policy problems, however, the envy-free efficient benchmark may be empty even after restricting attention to Pareto-efficient policies. Our goal is therefore to work with a selector that agrees with envy-free efficiency when it is attainable, and that yields a principled continuation when it is not.

Let N be a finite set of identities (agents, groups, or types) and let M be a finite set of states. A policy g induces, for each $i \in N$, a state-contingent consumption vector

$$c_i(g) \;=\; \big(c_{i,m}(g)\big)_{m\in M} \in \mathbb{R}^M,$$

and we write $c(g) = (c_i(g))_{i\in N}$ for the induced allocation profile. A policy menu G is a set of feasible policies generated by the underlying incentive, participation, budget, and institutional constraints. Policies are evaluated relative to a baseline policy g^0 (status quo or outside option), with $c_i^0 := c_i(g^0)$. The baseline is part of the economic problem: in the applications, reforms are naturally assessed as risky gains and losses relative to a risky status quo.

Preferences are denoted by $(\succeq_i)_{i\in N}$ on $\mathbb{R}^M$. In this section, we use only the Pareto order induced by these preferences; additional structure is imposed later when we translate risky changes into money-metric units.

Definition 1 (Pareto dominance and Pareto efficiency)**.** *Given a menu G, policy $g \in G$*

Pareto dominates $g' \in G$ *if* $c_i(g) \succeq_i c_i(g')$ *for all* $i \in N$ *and* $c_j(g) \succ_j c_j(g')$ *for some* $j \in N$. *A policy* $g \in G$ *is* Pareto efficient *in* G *if it is not Pareto dominated by any other policy in* G. *We write* $PE(G)$ *for the set of Pareto-efficient policies in* G.

In the environments we have in mind, $\mathrm{PE}(G)$ is the efficient locus generated by binding feasibility constraints and is often low-dimensional. Once efficiency is imposed, the fairness problem becomes a selection problem along this constrained frontier.

Our formal starting point is the surrogate approach developed in the companion paper *Fair choice under risk and uncertainty*. That paper takes as primitive a menu G together with an admissible evaluation profile $u^G = (u_i^G)_{i \in N}$. Each u_i^G is a within-menu cardinal index that is consistent with i's ranking on G and is interpreted as measuring i's gains in a common complaint unit along $\mathrm{PE}(G)$. The admissibility requirements are chosen so that the resulting selection depends only on complaint comparisons that are meaningful on a constrained efficient set, rather than on arbitrary normalizations of the indices.

Given an admissible pair (G, u^G), the companion paper assigns to each $g \in \mathrm{PE}(G)$ a residual-complaint (envy-slack) index $\mathrm{RES}(g; G, u^G)$, normalized so that larger values correspond to smaller residual complaints, and defines a selection correspondence

$$F(G; u^G) \subseteq \arg \max_{g \in \mathrm{PE}(G)} \mathrm{RES}(g; G, u^G),$$

together with a refinement that yields generic uniqueness and stability as menus vary within connected classes. Three implications are central for what follows. The selector always chooses Pareto-efficient policies. It coincides with envy-free efficiency whenever an envy-free efficient policy exists in the menu. When envy-free efficiency is unattainable, it provides a disciplined continuation on $\mathrm{PE}(G)$ by selecting frontier points that minimize the severity of residual complaints.[1]

For our purposes, $F(\cdot;\cdot)$ is the technical engine that turns within-menu complaint comparisons into an operational choice rule on constrained efficient frontiers. What remains is the measurement problem under risk: we need an evaluation profile u^G whose units are comparable across identities and economically interpretable in policy applications. Section 4 addresses this question by defining baseline-relative certainty-equivalent net gains and showing how they provide a canonical money-metric gauge for risky prospects.

4 Certainty-equivalent gauges and net-gain units

The surrogate selector in Section 3 compares Pareto-efficient policies by the severity of residual complaints. To use that selector in economics applications with risk, we must put complaints into a unit that is (i) preference grounded, so that it respects each identity's ranking of risky prospects, and (ii) money-metric, so that a sure dollar is valued at face value

[1] Continuation ideas appear in other feasibility-constrained settings. In matching with priorities, Abdulkadiroğlu et al. (2020) adapt the complaint logic by defining "justified envy" relative to priorities. With indivisibilities, approximate envy-freeness restores existence while preserving a complaint interpretation; see Plaut and Roughgarden (2020).

and incidence statements retain their usual meaning. Certainty equivalents provide precisely such a projection from risky prospects into dollars (Broome, 2012). Money-metric welfare analysis clarifies why a monetary gauge is the right object for interpersonal comparisons in applied settings with heterogeneous preferences (Blackorby and Donaldson, 1988; Fleurbaey, 2018).

As in Section 3, a policy g induces for each identity $i \in N$ a state-contingent consumption vector $c_i(g) \in \mathbb{R}^M$, and reforms are evaluated relative to a fixed baseline policy g^0 (status quo/outside option), writing $c_i^0 := c_i(g^0)$. In applications it is often convenient to keep deterministic payments (premia, taxes, subsidies) explicit. When we do so, we write $c_{i,m}(g) = y_{i,m}(g) - t_i(g)$ with $t_i(g) \in \mathbb{R}$ constant across states. This is purely bookkeeping: one can equivalently absorb $t_i(g)$ into $c_i(g)$ and work only with the induced consumption vector.

Assumption 1 (Expected-utility evaluation under (subjective) uncertainty)**.** *For each $i \in N$, there exist a strictly increasing and continuous Bernoulli utility function $v_i : \mathbb{R} \to \mathbb{R}$ and a probability vector $p_i \in \Delta(M)$ such that i evaluates $c \in \mathbb{R}^M$ by*

$$U_i(c) = \sum_{m \in M} p_{i,m}\, v_i(c_m).$$

When we use concavity properties of certainty equivalents, we additionally assume v_i is strictly concave.

Assumption 1 is used only to define a scalar certainty equivalent for each identity; it does not impose interpersonal comparability at the utility level. Heterogeneity in beliefs and risk attitudes enters through (p_i, v_i). The same idea extends to ambiguity-averse evaluation by replacing U_i with a robust aggregator and defining the associated robust certainty equivalent; Appendix A.2 develops this extension.

Definition 2 (Certainty equivalent)**.** *Under Assumption 1, the* certainty equivalent *of $c \in \mathbb{R}^M$ for identity i is the unique scalar $CE_i(c) \in \mathbb{R}$ satisfying*

$$v_i\big(CE_i(c)\big) = \sum_{m \in M} p_{i,m}\, v_i(c_m).$$

Definition 3 (CE gain and CE net gain)**.** *Fix the baseline g^0 and write $c_i^0 := c_i(g^0)$. The* CE gain *of g for i is*

$$\Delta_i(g) := CE_i\big(c_i(g)\big) - CE_i\big(c_i^0\big).$$

When policies include deterministic payments $t_i(g)$ that we wish to track separately, the CE net gain *is*

$$\Gamma_i(g) := \Delta_i(g) - t_i(g).$$

(Equivalently, one may absorb deterministic payments into $c_i(g)$ and work directly with $\Delta_i(g)$.)

The CE net gain $\Gamma_i(g)$ is the incidence object used throughout the paper. It is anchored at a common baseline, so it speaks directly to gains from reform. It is money-metric, so

sure transfers are treated as dollars and deterministic incidence terms can be separated without ambiguity. It is also preference grounded, because it evaluates risky changes using each identity's own beliefs and risk attitudes. Baseline anchoring is what makes this gauge economically meaningful in heterogeneous environments: it pins down the relevant normalization by measuring changes from a shared reference policy rather than from an arbitrary utility scale (Dávila and Schaab, 2025).

Two elementary properties of certainty equivalents are used repeatedly. First, they are monotone in the componentwise order, so a policy that improves consumption state by state weakly raises the certainty equivalent. Second, they respect the translation structure of money: adding a sure amount shifts the certainty equivalent by the same amount, and a sure outcome is valued at face value.

Lemma 1 (Basic properties of certainty equivalents)**.** *Suppose Assumption 1 holds. For each $i \in N$, the mapping $c \mapsto CE_i(c)$ is monotone, translation invariant with $CE_i(a\mathbf{1}) = a$ for all $a \in \mathbb{R}$, and continuous on $\mathbb{R}^M$. Consequently, $g \mapsto \Delta_i(g)$ and $g \mapsto \Gamma_i(g)$ are continuous whenever $g \mapsto c_i(g)$ and $g \mapsto t_i(g)$ are continuous.*

Example 2 (Closed form under CARA)**.** *If $v_i(x) = -\exp(-\alpha_i x)$ with $\alpha_i > 0$, then*

$$CE_i(c) = -\frac{1}{\alpha_i} \log\left(\sum_{m \in M} p_{i,m} \exp(-\alpha_i c_m) \right),$$

the entropic certainty equivalent. In the binary-loss insurance environments below, this yields explicit expressions for CE gains from coverage.

The next result clarifies in what sense certainty equivalents are canonical. Once we insist on a scalar index that preserves the identity's ranking over risky prospects and that treats sure transfers as additive dollars (with the normalization that a sure dollar is a dollar), the certainty equivalent is the unique such index.

Proposition 1 (Characterization of the certainty-equivalent gauge)**.** *Fix $i \in N$ and suppose Assumption 1 holds. Let $M_i : \mathbb{R}^M \to \mathbb{R}$ be any mapping satisfying:*

1. *(Preference consistency) $M_i(c) \geq M_i(c')$ if and only if $U_i(c) \geq U_i(c')$;*
2. *(Translation invariance) $M_i(c + a\mathbf{1}) = M_i(c) + a$ for all $c \in \mathbb{R}^M$ and $a \in \mathbb{R}$;*
3. *(Normalization on certainties) $M_i(a\mathbf{1}) = a$ for all $a \in \mathbb{R}$.*

Then $M_i(c) = CE_i(c)$ for all $c \in \mathbb{R}^M$.

Proof. Fix c and let $z := \mathrm{CE}_i(c)$. By definition of CE_i, $c \sim_i z\mathbf{1}$ under U_i. Preference consistency implies $M_i(c) = M_i(z\mathbf{1})$, and normalization yields $M_i(z\mathbf{1}) = z$. □

Proposition 1 is the identification step behind the informal instruction "measure risky gains in dollars." It also clarifies why the translation structure matters in constrained policy problems: if the scalar index does not treat sure dollars as additive dollars, then

incidence comparisons can change with irrelevant normalizations. For a related analysis of redistribution criteria based on certainty-equivalent comparisons, see Gauthier and Laroque (2024).

For applied work, CE net gains also admit a local representation in the same money units. Consider a one-parameter family of reforms $\{g^\varepsilon\}_\varepsilon$ with $g^0 = g^0$ and write $c^\varepsilon_{i,m} := c_{i,m}(g^\varepsilon)$ and $t^\varepsilon_i := t_i(g^\varepsilon)$. If v_i is differentiable and $\varepsilon \mapsto (c^\varepsilon_i, t^\varepsilon_i)$ is differentiable at $\varepsilon = 0$, differentiating the defining equation for $\mathrm{CE}_i(c^\varepsilon_i)$ yields

$$\frac{d}{d\varepsilon}\Gamma_i(g^\varepsilon)\Big|_{\varepsilon=0} = \sum_{m\in M} \omega^0_{i,m}\, \frac{d}{d\varepsilon} c^\varepsilon_{i,m}\Big|_{\varepsilon=0} - \frac{d}{d\varepsilon} t^\varepsilon_i\Big|_{\varepsilon=0}, \qquad \omega^0_{i,m} := \frac{p_{i,m}\, v_i'(c^0_{i,m})}{v_i'(\mathrm{CE}_i(c^0_i))}. \tag{1}$$

The weights $\omega^0_{i,m}$ tilt marginal valuations toward states with high marginal utility, so (1) expresses a marginal CE valuation as a risk-adjusted willingness-to-pay term minus a deterministic fiscal term. This is the same information base that underlies sufficient-statistics welfare analysis, but organized in identity-specific CE dollars rather than through externally chosen welfare weights (Finkelstein and Hendren, 2020; Hendren and Sprung-Keyser, 2020).

We now connect these CE units to the companion framework. The surrogate selector operates on a menu G together with a within-menu evaluation profile $u^G = (u^G_i)_{i\in N}$ that measures each identity's gains in a complaint unit along $\mathrm{PE}(G)$, up to normalizations that are immaterial for the selector. Baseline-relative CE net gains provide a canonical specialization of these indices: they place each identity's gain on a common money scale and thereby make residual-complaint comparisons economically interpretable on constrained efficient frontiers.

Proposition 2 (CE-indexed surrogate choice)**.** *For each menu G, define the evaluation profile u^G by $u^G_i(g) := \Gamma_i(g)$ for all $g \in G$ and $i \in N$. Under the regularity conditions of the companion paper for the pair (G, u^G), the resulting selection rule is Pareto efficient, coincides with the envy-free efficient benchmark when that benchmark is feasible, and otherwise implements the companion paper's minimal-complaint continuation on $PE(G)$ with residual complaints measured in CE net-gain dollars.*

This CE-indexed specialization is the object studied in the remainder of the paper. Section 5 characterizes the resulting fair-choice rule on constrained efficient frontiers and interprets it in terms of CE equalization and residual CE complaints.

5 Fair choice in CE net-gain units

Section 3 reduced fair policy choice under feasibility constraints to a selection problem on the Pareto frontier, and Section 4 fixed a canonical money-metric unit for comparing how identities fare under risk, the baseline-relative certainty-equivalent (CE) net gain $\Gamma_i(g)$. This section combines the two inputs. We specialize the companion paper's menu-based selector to the evaluation profile given by CE net gains and restate the resulting rule directly in incidence-relevant dollars. When the efficient frontier contains a policy that eliminates complaints in CE dollars, the rule implements that benchmark. When feasibility prevents

complaint elimination even on the frontier, the rule selects the efficient feasible policy that makes residual CE complaints as small as possible, with a leximin refinement when ties arise.

Fix a menu G and a baseline g^0. Define the CE-indexed rule by

$$F^{CE}(G) \;:=\; F(G; u^G), \qquad \text{where} \qquad u_i^G(g) := \Gamma_i(g) \;\text{ for all } i \in N,\, g \in G,$$

and $F(\cdot;\cdot)$ is the selector from the companion paper. Proposition 2 records the basic properties inherited from that framework, including Pareto efficiency and agreement with the envy-free efficient benchmark whenever that benchmark is feasible. The content of this section is to make the benchmark and its continuation explicit in CE net-gain units.

Definition 4 (Directed CE complaint and CE envy-freeness)**.** *For a policy $g \in G$ and identities $i, j \in N$, the* directed CE complaint *of i relative to j at g is*

$$e_{ij}^{CE}(g) \;:=\; \big(\Gamma_j(g) - \Gamma_i(g)\big)_+, \qquad (x)_+ := \max\{x, 0\}.$$

A policy g is CE envy-free *if $e_{ij}^{CE}(g) = 0$ for all $i, j \in N$.*

Directed CE complaints preserve the usual incidence interpretation. Identity i has a complaint against j precisely when j's CE net gain from the reform exceeds i's, and the complaint is measured in CE dollars. With finitely many identities, eliminating all directed complaints is equivalent to equalizing CE net gains.

Lemma 2 (CE envy-free $\iff$ CE equalization)**.** *A policy g is CE envy-free if and only if $\Gamma_i(g) = \Gamma_j(g)$ for all $i, j \in N$.*

Proof. If g is CE envy-free, then $e_{ij}^{CE}(g) = 0$ implies $\Gamma_i(g) \geq \Gamma_j(g)$ for all i, j, and symmetry yields equality. The converse is immediate. □

Let

$$\text{PEEF}^{\,CE}(G) \;:=\; \{g \in \text{PE}\,(G) : g \text{ is CE envy-free}\}$$

denote the CE envy-free efficient set. When $\text{PEEF}^{\,CE}(G)$ is nonempty, the benchmark is an egalitarian equalization in a money-metric unit. This is closely related in spirit to the egalitarian-equivalent comparisons of Pazner and Schmeidler (1978) and to money-metric approaches that separate the choice of a common metric from the ethical aggregation of welfare, as in d'Aspremont and Gevers (1977) and Fleurbaey and Maniquet (2011).

Proposition 3 (CE equalizer on the Pareto frontier)**.** *If $PEEF^{CE}(G) \neq \varnothing$, then every $g \in PEEF^{CE}(G)$ equalizes CE net gains across identities and the CE-indexed selector satisfies*

$$F^{CE}(G) \subseteq PEEF^{CE}(G).$$

In particular, under any refinement of the companion selector that is single-valued on $PEEF^{CE}(G)$, there is a unique CE equalizer *$g^{CE} \in PE(G)$ satisfying $\Gamma_i(g^{CE}) = \Gamma_j(g^{CE})$ for all i, j, and $F^{CE}(G) = \{g^{CE}\}$.*

Proposition 3 is the benchmark case: if the efficient locus permits full complaint elimination in CE dollars, fairness takes the form of equalizing risk-adjusted net gains measured in a common money-metric unit.

Feasibility constraints can prevent such equalization even on $\mathrm{PE}(G)$. In that case, the companion framework continues the benchmark by ranking Pareto-efficient policies according to the severity of their residual complaints. In CE units, the relevant residual object is the spread of CE net gains.

Definition 5 (Residual CE complaint)**.** *For a policy $g \in G$, define the* maximal CE gap *and the associated* residual CE complaint index *by*

$$D^{CE}(g) \;:=\; \max_{i,j\in N}\big(\Gamma_j(g)-\Gamma_i(g)\big) \;=\; \max_{j\in N}\Gamma_j(g)-\min_{i\in N}\Gamma_i(g), \qquad S^{CE}(g) \;:=\; -D^{CE}(g).$$

Thus $S^{CE}(g)\le 0$ for all g, and $S^{CE}(g)=0$ if and only if g is CE envy-free.

Theorem 1 (CE-indexed minimal-complaint continuation)**.** *Maintain the regularity conditions from the companion paper ensuring that maximizers exist on $PE(G)$, for example compactness of $PE(G)$ together with continuity of $g \mapsto \Gamma_i(g)$ as in Lemma 1. If $PEEF^{CE}(G)=\varnothing$, then the CE-indexed selector chooses Pareto-efficient policies that minimize the maximal CE gap, equivalently maximize S^{CE}, on the efficient frontier:*

$$F^{CE}(G) \;\subseteq\; \arg\min_{g\in PE(G)} D^{CE}(g) \;=\; \arg\max_{g\in PE(G)} S^{CE}(g).$$

Under the companion paper's refinement, which applies a leximin tie-break to the induced complaint vector, $F^{CE}(G)$ is generically single-valued.

Theorem 1 yields the practical take-away used throughout the applications: among efficient feasible reforms, choose the one that makes CE net gains as equal as possible, because that policy minimizes the worst residual complaint measured in CE dollars. This continuation logic mirrors two prominent ways the envy literature responds to infeasibility. Diamantaras and Thomson (1990) restore existence by relaxing no-envy parametrically, and Chaudhuri (1986) compare efficient outcomes by the severity of the envy they generate. The CE-indexed rule implements an intensity comparison on the efficient frontier, with intensity measured in a canonical money-metric unit under risk.

The continuation criterion admits an epigraph formulation that is useful both for computation and for interpretation. Appendix C shows that minimizing the maximal CE gap over the efficient frontier is equivalent to choosing (g,α,t) to minimize t subject to $g \in \mathrm{PE}(G)$ and

$$\alpha \le \Gamma_i(g) \le \alpha + t \qquad \text{for all } i \in N,$$

so that $t = D^{CE}(g)$ at the optimum (Proposition 8). The same proposition also provides a dual representation: for any fixed g, the gap $D^{CE}(g)$ equals the largest weighted difference in CE net gains between a distribution supported on the most-advantaged identities and a distribution supported on the most-disadvantaged identities. Thus the CE-indexed rule can be viewed as an efficient-policy choice that directly targets the CE-dollar spread between the extremes.

The two-group case used in the insurance application is especially transparent. Suppose there are two identities L, H and that efficient policies can be parameterized by a scalar $\theta \in [\underline{\theta}, \overline{\theta}]$ with $g(\theta) \in \mathrm{PE}(G)$ for all θ. Write $\Gamma_L(\theta) := \Gamma_L(g(\theta))$ and $\Gamma_H(\theta) := \Gamma_H(g(\theta))$.

Corollary 1 (Two groups: equalizer or closest-feasible equalization)**.** *Assume $\theta \mapsto \Gamma_L(\theta)$ and $\theta \mapsto \Gamma_H(\theta)$ are continuous on $[\underline{\theta}, \overline{\theta}]$ and that $\Gamma_L(\theta) - \Gamma_H(\theta)$ is strictly monotone in θ.*

1. *If there exists $\theta^\star \in [\underline{\theta}, \overline{\theta}]$ with $\Gamma_L(\theta^\star) = \Gamma_H(\theta^\star)$, then $\theta^\star$ is unique and any CE envy-free efficient policy equals $g(\theta^\star)$. Under a single-valued refinement, $F^{CE}(G) = \{g(\theta^\star)\}$.*

2. *If no such $\theta^\star$ exists, then every CE-indexed fair policy minimizes the absolute CE gap on the efficient frontier:*

$$F^{CE}(G) \;\subseteq\; \arg\min_{\theta \in [\underline{\theta}, \overline{\theta}]} \left|\Gamma_L(\theta) - \Gamma_H(\theta)\right|.$$

   If $\Gamma_L(\theta) - \Gamma_H(\theta)$ is strictly monotone, the minimizer is attained at a boundary point $\underline{\theta}$ or $\overline{\theta}$.

Proof. If $\Gamma_L(\theta) - \Gamma_H(\theta)$ crosses zero, strict monotonicity implies a unique root, yielding (i). Otherwise the absolute gap is minimized at the boundary point closest to zero. The link to $F^{CE}(G)$ follows from Theorem 1 and the identity $D^{CE}(g(\theta)) = \left|\Gamma_L(\theta) - \Gamma_H(\theta)\right|$ in the two-identity case. □

Section 6 shows that the same minimax and leximin logic has a veil-of-ignorance interpretation once regret is evaluated in CE net-gain dollars. Section 7 constructs the insurance menus, characterizes the corresponding efficient frontiers, derives explicit expressions for (Γ_L, Γ_H), and applies Corollary 1 to obtain closed-form characterizations and comparative statics.

6 Veil-of-ignorance foundations

Veil-of-ignorance arguments motivate fairness criteria by separating policy choice from identity: an evaluator chooses a feasible policy without knowing which identity she will occupy, and she anticipates assessing the realized incidence once identity is revealed. In our setting, uncertainty about physical outcomes has already been translated into a money-metric object through certainty equivalents (Section 4). The remaining uncertainty behind the veil concerns identity. This section shows that, when incidence is measured by baseline-relative CE net gains, minimizing worst-case *ex post* regret selects exactly the CE-indexed fair-choice rule from Section 5.

Fix a menu G and baseline g^0, and write $\Gamma(g) := (\Gamma_i(g))_{i \in N}$ for the profile of CE net gains induced by g. After choosing g, the evaluator learns which identity i she is and compares her own CE net gain to the most-advantaged CE net gain under the same chosen policy. This produces an *ex post* regret notion that is both preference grounded and money metric.

Definition 6 (Ex post CE envy and veil regret)**.** *For any policy $g \in G$ and identity $i \in N$, define the* ex post CE envy *(or regret) of identity i at g as*

$$R_i(g) \;:=\; \max_{j \in N}\big(\Gamma_j(g) - \Gamma_i(g)\big) \;=\; \max_{j \in N}\big(\Gamma_j(g) - \Gamma_i(g)\big)_+.$$

The associated worst-case veil regret *of g is*

$$R(g) \;:=\; \max_{i \in N} R_i(g).$$

This veil criterion is distinct from the two canonical uses of the veil. Harsanyi's impartial-observer argument treats the veil as uncertainty over utility and motivates ex ante aggregation (Harsanyi, 1955). Rawls uses the veil to motivate a maximin principle over social positions (Rawls, 1971). Here, the evaluator compares *baseline-relative incidence* in CE dollars and focuses on the largest ex post shortfall relative to the most-advantaged identity under the chosen policy. This emphasis on ex post comparisons connects to broader discussions of the tension between ex ante and ex post fairness under risk (Fleurbaey and Zuber, 2022) and to experimental evidence on veil reasoning (Andreoni et al., 2020).

A key simplification is that worst-case veil regret depends only on the spread of CE net gains.

Lemma 3 (Worst-case veil regret equals the maximal CE gap)**.** *For every $g \in G$,*

$$R(g) \;=\; \max_{i,j \in N}\big(\Gamma_j(g) - \Gamma_i(g)\big) \;=\; D^{CE}(g),$$

where $D^{CE}(g)$ is the maximal CE gap from Definition 5. In particular, $R(g) = 0$ if and only if g is CE envy-free.

Lemma 3 is an algebraic identity obtained by exchanging maxima. Appendix D records a general version and related auxiliary results.

We define veil choice by minimizing worst-case veil regret along the efficient frontier:

$$V^{CE}(G) \;:=\; \arg\min_{g \in \mathrm{PE}(G)} R(g). \tag{2}$$

When the minimizer is not unique, we apply the same refinement as in the companion paper, which coincides with the refinement invoked in Theorem 1.

Theorem 2 (Veil equivalence)**.** *Maintain the regularity conditions from Section 5 (and the companion paper) ensuring that the CE-indexed selector $F^{CE}(G)$ is well-defined and that minimizers exist on $PE(G)$. Then:*

1. *If $PEEF^{CE}(G) \neq \varnothing$, the veil criterion selects exactly the CE envy-free efficient set:*

$$V^{CE}(G) \;=\; PEEF^{CE}(G),$$

and, under a single-valued refinement, $V^{CE}(G) = F^{CE}(G) = \{g^{CE}\}$ where g^{CE} is the CE equalizer from Proposition 3.

2. *If* $PEEF^{CE}(G) = \varnothing$*, then the veil criterion coincides with the CE-indexed continuation problem:*

$$V^{CE}(G) \;=\; \arg\min_{g \in PE(G)} D^{CE}(g) \;=\; \arg\max_{g \in PE(G)} S^{CE}(g),$$

and, with the common refinement, $V^{CE}(G) = F^{CE}(G)$.

Proof. By Lemma 3, the veil problem is to minimize $D^{CE}(g)$ over $\mathrm{PE}(G)$. If $\mathrm{PEEF}^{CE}(G) \neq \varnothing$, the minimum value is 0, attained exactly at CE envy-free policies, which gives the first claim. If $\mathrm{PEEF}^{CE}(G) = \varnothing$, then 0 is unattainable and the veil problem is exactly the continuation program characterized in Theorem 1, which yields the second claim and the coincidence under the shared refinement. □

Theorem 2 provides the normative interpretation used in the remainder of the paper: the CE-indexed continuation rule is the policy choice of a veil-of-ignorance evaluator who anticipates judging the policy ex post, after identity is revealed, and who evaluates that shortfall in CE dollars. Related work develops veil interpretations of fairness criteria and studies how ex post concerns shape policy evaluation (Heufer et al., 2025; Liang and Lu, 2024b).

The two-group case that drives the insurance application reduces the veil objective to an absolute CE gap.

Corollary 2 (Two groups)**.** *Suppose there are two groups* L, H *and write* $\Gamma_L(g)$ *and* $\Gamma_H(g)$ *for their CE net gains. Then for any* $g \in G$*,*

$$R(g) \;=\; |\Gamma_L(g) - \Gamma_H(g)|.$$

Consequently,

$$V^{CE}(G) \;=\; \arg\min_{g \in PE(G)} |\Gamma_L(g) - \Gamma_H(g)|,$$

with equality whenever feasible (that is, whenever there exists $g \in PE(G)$ *with* $\Gamma_L(g) = \Gamma_H(g)$*). In particular, under the frontier parameterization of Corollary 1, the veil criterion selects the same frontier point as the CE-indexed surrogate choice.*

Proof. With two groups,

$$\max_{i,j}\big(\Gamma_j(g) - \Gamma_i(g)\big) = \max\{\Gamma_H(g) - \Gamma_L(g), \Gamma_L(g) - \Gamma_H(g)\} = |\Gamma_L(g) - \Gamma_H(g)|.$$

The remaining claims follow from Theorem 2 and Corollary 1. □

Section 7 applies this veil interpretation in an insurance environment where feasibility constraints restrict incidence along the efficient frontier. The resulting characterizations identify when CE equalization is feasible and, when it is not, which efficient feasible policy minimizes the worst ex post CE complaint.

7 Application: Fairness in insurance under adverse selection

This section applies the CE-indexed fair-choice rule to a canonical insurance environment in which feasibility constraints restrict both risk sharing and the incidence of financing. In the adverse-selection model of Rothschild and Stiglitz (1976), competitive contracting can unravel and block socially desirable pooling. Policy interventions respond through restricted instruments whose incidence is shaped by institutions, and welfare analysis in such environments is often framed in money-metric terms built from demand and cost primitives (Einav et al., 2010; Einav and Finkelstein, 2011). Our contribution here is different in scope: we treat the feasible instrument as part of the normative problem and show how the CE gauge and the continuation principle together yield an operational fairness prescription on the relevant efficient locus. The prescription equalizes CE net gains when the instrument can implement the required cross-subsidy, and otherwise selects the Pareto-efficient policy that makes the remaining CE gap as small as feasibility allows. Contribution ceilings provide a particularly transparent one-dimensional incidence instrument, and they are empirically important in social-insurance finance (Alvaredo et al., 2017).

There are two groups $i \in \{L, H\}$, interpreted as a disadvantaged high-risk and low-income group (L) and an advantaged low-risk and high-income group (H). Each group faces a binary loss of size $\ell > 0$, with primitives satisfying

$$q_L > q_H, \qquad 0 < q_i < 1, \qquad w_L < w_H. \tag{3}$$

Preferences and beliefs satisfy Assumption 1. For closed-form expressions we sometimes specialize to CARA utility,

$$v_i(c) = -\exp(-a_i c), \qquad a_i > 0, \tag{4}$$

but the fair-choice conclusions below rely only on continuity of CE net gains and on the one-dimensional incidence tradeoff induced by the financing constraint.

A policy specifies coverage $I \in [0, \ell]$ and deterministic payments $t_i \in \mathbb{R}$ for each group. Under policy $g = (I, t_L, t_H)$, consumption is

$$c_i(g) = \big(w_i - t_i,\ w_i - \ell + I - t_i\big),$$

where the first component is the no-loss state and the second the loss state. The baseline g^0 is no insurance, with $I = 0$ and $t_i = 0$, hence $c_i^0 = (w_i, w_i - \ell)$. Let $V_i(I)$ denote the CE value of coverage I relative to this baseline,

$$V_i(I) := \mathrm{CE}_i\big(w_i,\ w_i - \ell + I\big) - \mathrm{CE}_i\big(w_i,\ w_i - \ell\big), \tag{5}$$

so the CE net gain from policy g is

$$\Gamma_i(g) = V_i(I) - t_i. \tag{6}$$

Thus $V_i(\cdot)$ is the channel through which risk and risk preferences enter distributive incidence in CE dollars.

Lemma 4 (Value of coverage in CE units)**.** *Maintain Assumption 1. For each $i \in \{L, H\}$, the mapping $I \mapsto V_i(I)$ is continuous and strictly increasing on $[0, \ell]$. If v_i is strictly concave, then V_i is concave. Under CARA (4),*

$$V_i(I) = \frac{1}{a_i} \ln\Big((1 - q_i) + q_i e^{a_i \ell}\Big) - \frac{1}{a_i} \ln\Big((1 - q_i) + q_i e^{a_i(\ell - I)}\Big), \tag{7}$$

which is increasing in both q_i and a_i for each fixed $I > 0$.

We first study mandated public insurance. Absent moral hazard, full insurance $I = \ell$ is efficient, so the policy problem is incidence: who finances the pooled coverage. We model a restricted financing instrument by a proportional contribution rate applied to an earnings base capped at $\bar{w}$ (a contribution ceiling). For $\bar{w} \in [w_L, w_H]$, define

$$t_i(\bar{w}) = \tau(\bar{w}) \min\{w_i, \bar{w}\}, \qquad \tau(\bar{w}) = \frac{(q_L + q_H)\ell}{\min\{w_L, \bar{w}\} + \min\{w_H, \bar{w}\}}, \tag{8}$$

so that budget balance under full insurance holds:

$$t_L(\bar{w}) + t_H(\bar{w}) = (q_L + q_H)\ell. \tag{9}$$

Given (3) and $\bar{w} \in [w_L, w_H]$, we have $\min\{w_L, \bar{w}\} = w_L$ and $\min\{w_H, \bar{w}\} = \bar{w}$, hence

$$t_L(\bar{w}) = \frac{(q_L + q_H)\ell\, w_L}{w_L + \bar{w}}, \qquad t_H(\bar{w}) = \frac{(q_L + q_H)\ell\, \bar{w}}{w_L + \bar{w}}. \tag{10}$$

Let $g^{\text{man}}(\bar{w})$ denote the mandated full-insurance policy with ceiling $\bar{w}$,

$$g^{\text{man}}(\bar{w}) = \big(\ell, t_L(\bar{w}), t_H(\bar{w})\big),$$

and let $G^{\text{man}} = \{g^{\text{man}}(\bar{w}) : \bar{w} \in [w_L, w_H]\}$ denote the mandated menu. Along this menu, the ceiling shifts incidence through (t_L, t_H) while keeping coverage fixed at the efficient level.

The two-group characterization from Section 5 applies because the ceiling induces a monotone tradeoff in the CE net-gain difference.

Lemma 5 (Monotone incidence along the mandated frontier)**.** *Along G^{man}, the CE net-gain difference $\bar{w} \mapsto \Gamma_L(g^{\text{man}}(\bar{w})) - \Gamma_H(g^{\text{man}}(\bar{w}))$ is strictly increasing on $[w_L, w_H]$.*

The next theorem states the incidence content of the CE-indexed rule. Since the ceiling is the only degree of freedom on the efficient locus, the fair policy is characterized by whether that single instrument can implement CE equalization.

Theorem 3 (Fair incidence under mandated social insurance)**.** *Maintain (3). Consider the mandated menu G^{man} and apply the CE-indexed selector F^{CE} (Section 5).*

1. *If there exists $\bar{w}^* \in [w_L, w_H]$ such that*

$$\Gamma_L\big(g^{\text{man}}(\bar{w}^*)\big) = \Gamma_H\big(g^{\text{man}}(\bar{w}^*)\big), \tag{11}$$

then $\bar{w}^*$ is unique and the CE-indexed fair policy is $g^{\text{man}}(\bar{w}^*)$.

2. *If no $\bar{w}^*$ satisfies (11), then CE equalization is infeasible on G^{man} and the CE-indexed fair policy is the boundary ceiling (either w_L or w_H) that minimizes the absolute CE net-gain gap*

$$\left|\Gamma_L\left(g^{\text{man}}(\bar{w})\right) - \Gamma_H\left(g^{\text{man}}(\bar{w})\right)\right|.$$

Equivalently, it minimizes worst-case veil regret in CE dollars among Pareto-efficient mandated policies.

Under full insurance, $\Gamma_i = V_i(\ell) - t_i(\bar{w})$, so CE equalization is equivalent to the incidence condition

$$t_H(\bar{w}) - t_L(\bar{w}) = V_L(\ell) - V_H(\ell).$$

The left-hand side is pinned down by the ceiling via (10), while the right-hand side is the CE valuation gap induced by risk and risk preferences through $V_i(\ell)$. When the statutory range of $\bar{w}$ is rich enough, the fair ceiling is the unique point on the efficient locus that matches these two objects. When institutions restrict $\bar{w}$ so that the condition is unattainable, the rule continues by selecting the boundary instrument that gets closest, and the minimal achievable CE gap provides a money-metric measure of how tightly feasibility constrains distributive incidence.

We now allow enrollment to be voluntary. A pooled policy specifies coverage I and a ceiling $\bar{w}$, with balanced-budget contributions given by

$$t_i(I, \bar{w}) = \tau(I, \bar{w}) \min\{w_i, \bar{w}\}, \qquad \tau(I, \bar{w}) = \frac{(q_L + q_H)I}{\min\{w_L, \bar{w}\} + \min\{w_H, \bar{w}\}}. \tag{12}$$

Voluntary participation requires that each group weakly prefers enrollment to the baseline,

$$\Gamma_i(I, \bar{w}) = V_i(I) - t_i(I, \bar{w}) \geq 0, \qquad i \in \{L, H\}. \tag{13}$$

Let G^{pool} denote the set of pooled policies satisfying (12)–(13). When G^{pool} is empty, pooling cannot be sustained and adverse selection implies separation in the usual sense of Rothschild and Stiglitz (1976).

We represent separation by a benchmark in which group L (high risk) can purchase actuarially fair coverage I at premium $q_L I$, while group H opts out and remains at the baseline. Thus a separating policy is $g^{\text{sep}}(I)$ with

$$\Gamma_L\left(g^{\text{sep}}(I)\right) = V_L(I) - q_L I, \qquad \Gamma_H\left(g^{\text{sep}}(I)\right) = 0. \tag{14}$$

Absent moral hazard, $I = \ell$ is efficient among separating policies because it maximizes Γ_L.

The next theorem expresses pooling versus separation in CE dollars. When pooling is feasible, the problem is again selection along an efficient constrained frontier, now shaped jointly by budget balance and participation constraints. When pooling is infeasible, the separating regime is efficient but forces a strictly negative residual slack that quantifies the distributive wedge created by adverse selection within the restricted menu.

Theorem 4 (Fair policy under voluntary enrollment)**.** *Maintain* (3).

1. *If the pooled set G^{pool} is nonempty, then the CE-indexed fair policy is a Pareto-efficient pooled policy that minimizes the absolute CE net-gain gap on the pooling Pareto frontier:*

$$F^{CE}(G^{\text{pool}}) \subseteq \arg \min_{g \in PE(G^{\text{pool}})} |\Gamma_L(g) - \Gamma_H(g)|.$$

   In particular, when the pooling efficient locus is one-dimensional with a monotone incidence tradeoff in $\Gamma_L - \Gamma_H$, the fair pooled policy equalizes CE net gains whenever feasible and otherwise lies at the boundary point that makes the gap as small as feasibility allows.

2. *If G^{pool} is empty, pooling is infeasible and the efficient separating regime yields $\Gamma_H = 0$ and $\Gamma_L = V_L(\ell) - q_L\ell > 0$. The resulting residual CE slack is*

$$S^{CE,\text{sep}} = -|\Gamma_L - \Gamma_H| = -(V_L(\ell) - q_L\ell) < 0,$$

   which provides a money-metric measure of the fairness loss forced by adverse selection.

The separating slack $|S^{CE,\text{sep}}| = V_L(\ell) - q_L\ell$ isolates the part of the pooling-versus-separation tension that is purely distributive within the constrained menu: it is the minimal CE-dollar spread that feasibility forces once pooling is ruled out. This complements the standard efficiency accounting in the empirical selection literature. Einav and Finkelstein (2011) provide the central demand-and-cost welfare formulas, Einav and Finkelstein (2023) survey the modern methods, and Rothschild (2024) refines welfare accounting under selection. The same CE-dollar slack is also comparable across regulated environments in which mandates and plan design shift the incidence of selection. For example, Hackmann et al. (2015) quantify incidence effects of mandates in regulated markets, while recent work studies how networks and other design restrictions reshape incidence under regulation (Kreider et al., 2024; Kong et al., 2024).

Two take-home messages summarize the application. First, under mandated insurance with a contribution ceiling, the efficient locus is one-dimensional and the fair ceiling equalizes CE net gains when feasible; otherwise it implements the closest-feasible CE equalization on the Pareto frontier (Theorem 3). Second, under voluntary enrollment, if pooling is feasible the fair pooled policy is the Pareto-efficient pooled policy that makes the CE net-gain gap as small as feasibility allows; if pooling is infeasible, separation forces a strictly negative residual slack $S^{CE,\text{sep}} = -(V_L(\ell) - q_L\ell)$ that summarizes the fairness loss from unraveling (Theorem 4).

The CE formulation also makes comparative statics immediate. Changes in risk or risk preferences shift willingness to pay for coverage through $V_i(\cdot)$ (Lemma 4). Under CARA (4), higher q_L or a_L increases $V_L(\ell)$ and therefore raises the incidence transfer required for CE equalization. When equalization is feasible, this pushes the fair mandated ceiling toward greater cross-subsidization of group L, because a higher ceiling increases $t_H(\bar{w}) - t_L(\bar{w})$ in (10). Tightening the financing instrument, for example by restricting the admissible range of $\bar{w}$, can make equalization infeasible. In that case, the fair policy moves to the closest feasible

boundary and the minimal achievable gap $|\Gamma_L - \Gamma_H|$ becomes a money-metric measure of the distributive distortion induced by constraints.

This section abstracts from moral hazard in order to isolate incidence and adverse-selection constraints. With moral hazard, coverage I becomes a genuine choice variable and the efficient locus is generally multidimensional, but the CE-indexed criterion continues to apply directly through the continuation programs in Section 5. Screening models with incentive-compatible contract menus can be incorporated by letting G represent the set of incentive-feasible policies and evaluating each policy by its induced CE net gains. Appendix E provides the frontier characterizations invoked above and develops these extensions.

8 Wealth taxation

Wealth-tax reform illustrates a constrained-menu problem in which statutory form and administrative capacity shape incidence even before behavioral responses are introduced. Any implementable reform must respect a legally admissible functional form (brackets, base definitions, valuation rules, enforcement technology, and feasible rebate schemes). At the same time, risk is central: households face uncertainty in labor income and, for the very wealthy, in asset returns and valuations. This makes baseline-relative certainty-equivalent (CE) net gains a natural incidence unit, because they aggregate mean disposable-income changes and changes in risk exposure in a single money metric.

A useful benchmark is the classic optimal-tax tradition, which emphasizes the joint design of redistribution and incentives. Mirrlees (1971) introduces the optimal income-tax problem, Atkinson and Stiglitz (1976) characterize when commodity taxation is redundant, Saez (2001) derives implementable optimal-tax formulas, and Saez and Stantcheva (2016) organize welfare analysis using social marginal welfare weights. The wealth-tax debate emphasizes a complementary set of constraints: the effectiveness and incidence of wealth taxation depend on base definition, valuation, enforcement, and avoidance opportunities, as synthesized by Scheuer and Slemrod (2021) and illustrated empirically by Jakobsen et al. (2020). Our framework treats these statutory and administrative features as defining the implementable menu and then evaluates candidate reforms on the resulting constrained efficient set.

To map a statutory proposal into our model, take the status quo tax system as the baseline policy g^0. Let $y_{i,m}$ denote identity i's baseline after-transfer resources in state $m \in M$, and consider a restricted menu of wealth-tax-and-rebate reforms indexed by an implementable parameter $\theta \in [\underline{\theta}, \overline{\theta}]$. A reform $g(\theta)$ specifies (i) a progressive annual wealth-tax schedule $T_\theta(\cdot)$ applied to end-of-year net worth and (ii) a deterministic rebate rule that returns revenue as a (possibly targeted) lump-sum transfer. In the general model, the wealth tax alters the state-contingent consumption vector $c_i(g(\theta))$ through realized wealth and returns, while the rebate enters as a state-invariant component of net transfers.[2]

[2]Equivalently, one may incorporate the rebate into c_i as a state-invariant shift; by translation invariance of certainty equivalents, either representation yields the same Γ_i, cf. Definition 3.

A salient statutory form is the progressive wealth-tax proposal associated with Saez and Zucman (2019a) and Saez and Zucman (2019b). Abstracting from definitional and enforcement details, the schedule can be written as

$$T^{SZ}(W) = 0.02\,(W - 50\text{m})_+ + 0.01\,(W - 1\text{b})_+,$$

so that the marginal rate is 2% above \$50 million and 3% above \$1 billion. To preserve the statutory bracket structure while allowing implementable variation, it is convenient to index the family as $T_\theta(W) = \theta\, T^{SZ}(W)$, where θ captures rate scaling, effective coverage/enforcement, or administratively admissible adjustments that shift incidence without changing the schedule's form.[3]

A numerical example helps fix magnitudes and avoids conflating marginal and average rates. Under the statutory proposal ($\theta = 1$), a household with net worth $W = \$100$ million owes $T^{SZ}(W) = 0.02 \times (\$50\text{m}) = \$1$ million, an average rate of 1%.[4] A household with $W = \$200$ million owes $T^{SZ}(W) = 0.02 \times (\$150\text{m}) = \$3$ million, an average rate of 1.5%. A household with $W = \$2$ billion owes

$$T^{SZ}(W) = 0.02 \times (\$1.95\text{b}) + 0.01 \times (\$1\text{b}) = \$49\text{m},$$

an average rate of 2.45% of total wealth. Under $T_\theta(W) = \theta\, T^{SZ}(W)$, each of these liabilities and average rates scales linearly with θ.

For our purposes, the key economic content is that wealth-tax liabilities are state contingent for the liable identity: when wealth is risky, tax payments are higher in high-return or high-valuation states and lower in low states. As a result, the liable identity's CE net gain need not move one-for-one with expected statutory payments, because the reform also changes the distribution of resources across states and therefore the insurance content of after-tax wealth. Conversely, a non-liable identity can receive a relatively safe rebate financed by a relatively risky tax base, generating a sharp CE-incidence tradeoff even absent behavioral responses. This is precisely the setting in which comparisons based solely on expected disposable-income changes can be misleading, while CE net gains remain interpretable.

To connect to the two-identity selector from Section 5, let L denote an identity that is not directly liable for the wealth tax (or has negligible taxable wealth relative to the exemption) and let H denote an identity that is liable. Suppose the restricted efficient locus within the statutory class can be parameterized by $\theta \mapsto g(\theta) \in \mathrm{PE}\,(G)$ and that θ shifts incidence monotonically in the sense that higher θ raises L's CE net gain and lowers H's CE net gain. Then $\theta \mapsto \Gamma_L(g(\theta)) - \Gamma_H(g(\theta))$ is monotone, and Corollary 1 yields an operational fair-choice prescription within the statutory class: if there exists $\theta^\star$ with $\Gamma_L(g(\theta^\star)) = \Gamma_H(g(\theta^\star))$, the CE-indexed rule selects $g(\theta^\star)$; otherwise it selects the boundary parameter that minimizes $|\Gamma_L(g(\theta)) - \Gamma_H(g(\theta))|$, equivalently minimizing worst-case veil regret in CE dollars. Appendix Proposition 18 formalizes this evaluation and defines a money-metric fairness shortfall for

[3] The proposal discusses avoidance and valuation issues and makes explicit enforcement assumptions. In our framework these features are represented as feasibility restrictions that constrain the implementable set G, and hence the admissible range $[\underline{\theta}, \overline{\theta}]$ and the mapping $\theta \mapsto c_i(g(\theta))$.

[4] This calculation (and the analogous example at \$2 billion below) appears in Saez and Zucman (2019b).

adopting the statutory proposal within the class, while Appendix Proposition 19 provides a general optimality certificate for whether a given statutory proposal already minimizes the maximal CE complaint within the implementable parameter set.

Empirically, implementation requires the same core inputs as distributional tax analysis, but organized in CE units. Given an identity partition (for example, wealth-tax-liable versus non-liable, or top 0.1% versus the rest), one can simulate state-contingent after-tax resources under $g(\theta)$ using observed wealth composition, return processes, valuation risk, and enforcement/avoidance assumptions, and then compute $\Gamma_i(g(\theta))$ directly. For marginal changes within the statutory family, the local representation in (1) expresses $d\Gamma_i/d\theta$ as a risk-adjusted willingness-to-pay term minus deterministic fiscal consequences. The resulting CE-incidence frontier can then be compared across alternative wealth-tax designs and enforcement regimes, complementing quantitative analyses that emphasize behavioral elasticities and long-run accumulation responses (Jakobsen et al., 2020; Guvenen et al., 2023).

9 Further applications and extensions

The insurance application highlights a general structure that recurs in policy work. Feasibility constraints define an implementable menu, Pareto efficiency restricts attention to a constrained frontier, and the remaining distributive choice is often organized by a small number of incidence instruments. Once incidence is measured in baseline-relative certainty-equivalent (CE) net gains, the CE-indexed rule from Sections 4–6 yields a direct prescription on that frontier. This section records a reusable two-group template and sketches two additional environments in which it applies.

Many applications reduce to two identities and a one-dimensional efficient locus. The next assumption collects the regularity conditions under which the two-group characterization from Section 5 becomes a plug-in.

Assumption 2 (Two-group frontier regularity)**.** *There are two identities L, H and a feasible menu G with Pareto frontier $PE(G)$. There exists a parameterization $\theta \mapsto g(\theta) \in PE(G)$ for $\theta \in [\underline{\theta}, \overline{\theta}]$ such that the induced CE net gains $\Gamma_L(\theta) := \Gamma_L(g(\theta))$ and $\Gamma_H(\theta) := \Gamma_H(g(\theta))$ are continuous and $\Gamma_L(\theta) - \Gamma_H(\theta)$ is strictly monotone in θ.*

Assumption 2 is tailored to settings in which θ is an incidence instrument, a financing rule, eligibility threshold, or statutory parameter, that shifts who pays and who benefits while the underlying risk technology is held fixed. In such cases, strict monotonicity of $\Gamma_L - \Gamma_H$ is a transparent CE-incidence statement.

Proposition 4 (A reusable two-group template)**.** *Under Assumption 2, the CE-indexed fair choice and its veil interpretation are characterized by Corollary 1 and Corollary 2. If there exists $\theta^\star$ with $\Gamma_L(\theta^\star) = \Gamma_H(\theta^\star)$, the unique fair policy is $g(\theta^\star)$. If no such $\theta^\star$ exists, the fair policy is the boundary point $\underline{\theta}$ or $\overline{\theta}$ that minimizes $|\Gamma_L(\theta) - \Gamma_H(\theta)|$.*

Proposition 4 reduces an application to three economic tasks: specify the menu G by

making feasibility constraints explicit; characterize the efficient locus and isolate an incidence parameter θ; and verify that θ shifts the CE gap monotonically.

Unemployment insurance and wage insurance provide a natural illustration. Shavell and Weiss (1979) study unemployment insurance design under moral hazard, and Hopenhayn and Nicolini (1997) develop dynamic contracting formulations of the same efficiency problem. In practice, reforms are implemented through restricted instruments such as replacement-rate formulas, benefit duration, eligibility rules, and payroll-tax financing caps, each of which generates a constrained menu and typically a low-dimensional incidence tradeoff on the efficient set. CE net gains are useful in these environments because they aggregate expected-resource changes and changes in risk exposure in one money-metric object, which is exactly what is needed when liquidity and insurance value are first-order (Chetty, 2008). Recent empirical work similarly emphasizes that reforms can shift both behavior and the value of insurance, making a risk-adjusted incidence unit central for normative comparisons (Wunsch and Zabrodina, 2023). Wage insurance proposals operate through related statutory instruments (eligibility thresholds and subsidy rates), and their incidence is often governed by a small number of parameters in the implementable class; see Hyman et al. (2024). When the efficient locus within the admissible instrument class can be indexed by a single incidence parameter, Proposition 4 identifies the fair reform by CE equalization when feasible and otherwise by closest-feasible equalization in CE dollars.

Credit policy offers a second illustration. With adverse selection, moral hazard, and collateral constraints, implementable lending and repayment schedules may generate rationing rather than price adjustment, as in Stiglitz and Weiss (1981). Policy interventions such as loan guarantees, subsidized credit, or collateral relief therefore operate through a constrained menu and change both access and fiscal incidence. Taking the prevailing regime as the baseline, CE net gains summarize how a reform changes borrowers' consumption risk through default exposure and repayment obligations while capturing the deterministic incidence borne by others through taxes or fees. In large-scale guarantee programs, a small number of statutory parameters often organize the central access–incidence tradeoff (Altavilla et al., 2021), and the same parameters shape how risk is allocated across borrowers, lenders, and the fiscal authority (Carletti et al., 2023). When this structure yields a one-dimensional efficient locus, the fair choice again follows from Proposition 4.

Several extensions follow directly from the general formulation. When moral hazard or contract design makes coverage and incentives genuine choice variables, Pareto-efficient policies may form a multidimensional set; the CE-indexed criterion continues to apply via the continuation programs in Section 5, which minimize complaints in CE dollars on $\mathrm{PE}(G)$. With more than two identities or continuous heterogeneity, the same logic applies using the general max–min and leximin refinements, without relying on the two-group reduction. Ambiguity about physical states can be incorporated by replacing expected-utility certainty equivalents with robust certainty equivalents, as developed in Appendix A.2. Finally, because CE net gains are baseline-relative, the same selection logic applies naturally to reforms evaluated against institutionally salient status quo policies, which is the relevant comparison in policy environments.

10 Discussion

This paper combines two ingredients that are often bundled together in applied evaluation under risk. The first is an interpersonal *unit* for comparing policy consequences. We measure each identity's consequences in baseline-relative certainty-equivalent (CE) net gains (Section 4). This is a money-metric projection of risky prospects that treats sure dollars as dollars and evaluates risk using each identity's own beliefs and risk attitudes (Broome, 2012). Money-metric welfare analysis clarifies why fixing such a monetary gauge is the right starting point when preferences differ across people (Blackorby and Donaldson, 1988). The second ingredient is a *selection rule* for choosing from a constrained menu of feasible reforms. We take feasibility constraints as primitive and restrict attention to Pareto-efficient policies (Section 3). Conditional on the CE gauge, the selector is complaint based: it agrees with envy-free efficiency when attainable and remains well-defined when it is not (Section 5). This complaint orientation follows the envy-free tradition (Foley, 1967; Varian, 1974) while adapting it to feasibility-constrained policy menus under risk.

A familiar alternative in public finance is to aggregate money-metric gains using welfare weights. With population shares, or an impartial-observer prior, $\mu \in \Delta(N)$, the utilitarian benchmark in CE dollars is

$$W^U_\mu(g) \;=\; \sum_{i \in N} \mu_i \, \Gamma_i(g), \tag{15}$$

where $\Gamma_i(g)$ is the CE net gain relative to the baseline (Definition 3). This corresponds to a Harsanyi-style veil once outcomes have been expressed in CE dollars (Harsanyi, 1955). In applied work, the same weighted-sum language underlies sufficient-statistics and MVPF-style evaluation frameworks (Saez and Stantcheva, 2016; Hendren and Sprung-Keyser, 2020; Finkelstein and Hendren, 2020). Prioritarian criteria apply a concave transformation before summing,

$$W^P_\mu(g) \;=\; \sum_{i \in N} \mu_i \, \phi\big(\Gamma_i(g)\big), \tag{16}$$

with ϕ increasing and concave, which builds in smooth inequality aversion in the same units.

Our rule differs in what it asks the planner to commit to. It is not intended as a single welfare function that ranks all policies across all environments. Instead it is a menu selector disciplined by requirements that become salient when feasibility constraints are central. The rule (i) selects only Pareto-efficient policies, (ii) coincides with envy-free efficiency whenever that benchmark is feasible on the efficient locus, and (iii) when equalization is infeasible, continues intrinsically along the constrained frontier by minimizing the maximal residual complaint in CE dollars (Theorem 1). The continuation is closely related to two ways the envy literature responds to infeasibility: relaxing no-envy to restore existence (Diamantaras and Thomson, 1990) and comparing efficient allocations by the severity of envy they generate (Chaudhuri, 1986). The present paper's contribution is to implement this intensity comparison in a canonical money-metric unit under risk, so that the residual complaint has an immediate incidence interpretation.

The veil-of-ignorance result in Section 6 clarifies the normative content of this continuation. A Harsanyi veil evaluates policies ex ante via a weighted sum (as in (15)) (Harsanyi,

1955), while Rawlsian reasoning emphasizes protecting the worst position behind the veil (Rawls, 1971). Our equivalence result identifies the CE-indexed continuation with minimizing worst-case *ex post* regret in CE net-gain dollars on the Pareto frontier (Theorem 2). This criterion is informative precisely when feasibility prevents equalization, and it makes the feasibility dependence explicit by restricting evaluation to $\mathrm{PE}(G)$. The result also connects to broader discussions of tensions between ex ante and ex post principles under risk and heterogeneous preferences (Fleurbaey and Zuber, 2022; Miyagishima, 2022; Fleurbaey and Zuber, 2025).

The CE gauge makes heterogeneity economically transparent. Each $\Gamma_i(g)$ is computed using identity i's own Bernoulli utility and beliefs (Assumption 1), so disagreement about risk exposure or risk preferences appears as disagreement about willingness to pay for risk reduction, stated in dollars. The fairness comparison then asks whether the reform distributes these risk-adjusted gains relative to the status quo fairly across identities, rather than comparing expected-utility levels across people. The local representation in Section 4 connects Γ_i to empirically accessible objects, which helps bridge normative evaluation and sufficient-statistics implementation without requiring a commitment to a global welfare-weight schedule.

Several directions follow directly from the "menu first" discipline. Many policy environments generate multidimensional efficient sets, for example with moral hazard, screening, or multi-instrument regulation; in those cases the continuation programs in Section 5 still provide an operational criterion on $\mathrm{PE}(G)$, and Appendix C gives computational representations. Dynamic environments add intertemporal feasibility constraints and evolving risks, strengthening the case for treating the implementable menu as primitive. Finally, statutory reforms, including wealth taxation, often hinge on enforcement and avoidance constraints that define the implementable class; the optimality certificates in Subsection 8 are designed to separate disagreements about feasibility from disagreements about distributive objectives.

References

Abdulkadiroğlu, A., Che, Y.-K., Pathak, P. A., Roth, A. E., and Tercieux, O. (2020). Efficiency, justified envy, and incentives in priority-based matching. *American Economic Review: Insights*, 2(4):425–442.

Akerlof, G. A. (1970). The market for "lemons": Quality uncertainty and the market mechanism. *The Quarterly Journal of Economics*, 84(3):488–500.

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

Altavilla, C., Ellul, A., Pagano, M., Polo, A., and Vlassopoulos, T. (2021). Loan guarantees, bank lending and credit risk reallocation. Working Paper 629, CSEF (University of Naples Federico II).

Alvaredo, F., Breda, T., Roantree, B., and Saez, E. (2017). Contribution ceilings and the incidence of payroll taxes. *De Economist*, 165:129–140.

Andreoni, J., Aydin, D., Barton, B., Bernheim, B. D., and Naecker, J. (2020). When fair isn't fair: Understanding choice reversals involving social preferences. *Journal of Political Economy*, 128(5):1673–1711.

Atkinson, A. B. and Stiglitz, J. E. (1976). The design of tax structure: Direct versus indirect taxation. *Journal of Public Economics*, 6(1-2):55–75.

Blackorby, C. and Donaldson, D. (1988). Money metric utility: A harmless normalization? *Journal of Economic Theory*, 46(1):120–129.

Bogomolnaia, A. and Moulin, H. (2001). A new solution to the random assignment problem. *Journal of Economic Theory*, 100(2):295–328.

Broome, J. (2012). *Climate Matters: Ethics in a Warming World*. W. W. Norton & Company.

Budish, E. (2011). The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy*, 119(6):1061–1103.

Carletti, E., Leonello, A., and Marquez, R. (2023). Loan guarantees, bank underwriting policies and financial fragility. Working Paper Series 2782, European Central Bank.

Chaudhuri, A. (1986). Some implications of an intensity measure of envy. *Social Choice and Welfare*, 3(4):255–270.

Chetty, R. (2008). Moral hazard versus liquidity and optimal unemployment insurance. *Journal of Political Economy*, 116(2):173–234.

d'Aspremont, C. and Gevers, L. (1977). Equity and the informational basis of collective choice. *Review of Economic Studies*, 44(2):199–209.

Dávila, E. and Schaab, A. (2025). Welfare assessments with heterogeneous individuals. *Journal of Political Economy*, 133(9):2918–2961.

Diamantaras, D. and Thomson, W. (1990). A refinement and extension of the no-envy concept. *Economics Letters*, 33(3):217–222.

Einav, L. and Finkelstein, A. (2011). Selection in insurance markets: Theory and empirics in pictures. *Journal of Economic Perspectives*, 25(1):115–138.

Einav, L. and Finkelstein, A. (2023). Empirical analyses of selection and welfare in insurance markets: A self-indulgent survey. *The Geneva Risk and Insurance Review*, 48(2):167–191.

Einav, L., Finkelstein, A., and Cullen, M. R. (2010). Estimating welfare in insurance markets using variation in prices. *The Quarterly Journal of Economics*, 125(3):877–921.

Finkelstein, A. and Hendren, N. (2020). Welfare analysis meets causal inference. *Journal of Economic Perspectives*, 34(4):146–167.

Fleurbaey, M. (2009). Beyond gdp: The quest for a measure of social welfare. *Journal of Economic literature*, 47(4):1029–75.

Fleurbaey, M. (2018). Welfare economics, risk and uncertainty. *Canadian Journal of Economics*, 51(1):5–40.

Fleurbaey, M. and Maniquet, F. (2011). *A Theory of Fairness and Social Welfare*. Cambridge University Press.

Fleurbaey, M. and Zuber, S. (2022). A new puzzle in the social evaluation of risk. *Economics & Philosophy*, 38(3):450–465.

Fleurbaey, M. and Zuber, S. (2025). Universal social welfare orderings and risk. *Journal of Economic Theory*, 226:106014.

Foley, D. K. (1967). Resource allocation and the public sector. *Yale Economic Essays*, 7(1):45–98.

Gauthier, S. and Laroque, G. (2024). Weak redistribution and certainty equivalent domination. *Journal of Public Economic Theory*, 26(3):e12703.

Guvenen, F., Kambourov, G., Kuruscu, B., Ocampo, S., and Chen, D. (2023). Use it or lose it: Efficiency and redistributional effects of wealth taxation. *The Quarterly Journal of Economics*, 138(2):835–894.

Hackmann, M. B., Kolstad, J. T., and Kowalski, A. E. (2015). Adverse selection and an individual mandate: When theory meets practice. *American Economic Review*, 105(3):1030–1066.

Harsanyi, J. C. (1955). Cardinal welfare, individualistic ethics, and interpersonal comparisons of utility. *Journal of Political Economy*, 63(4):309–321.

Hendren, N. and Sprung-Keyser, B. (2020). A unified welfare analysis of government policies. *The Quarterly Journal of Economics*, 135(3):1209–1318.

Heufer, J., Shachat, J., and Xu, Y. (2025). Measuring tastes for equity and aggregate wealth behind the veil of ignorance. *Games and Economic Behavior*, 153:209–232.

Hopenhayn, H. A. and Nicolini, J. P. (1997). Optimal unemployment insurance. *Journal of Political Economy*, 105(2):412–438.

Hyman, B. G., Kovak, B. K., and Leive, A. (2024). Wage insurance for displaced workers. NBER Working Paper 32464, National Bureau of Economic Research.

Jakobsen, K., Jakobsen, K., Kleven, H., and Zucman, G. (2020). Wealth taxation and wealth accumulation: Theory and evidence from denmark. *The Quarterly Journal of Economics*, 135(1):329–388.

Kong, E., Layton, T. J., and Shepard, M. (2024). Adverse selection and (un)natural monopoly in insurance markets. NBER Working Paper 33187, National Bureau of Economic Research.

Kreider, A. R., Layton, T. J., Shepard, M., and Wallace, J. (2024). Adverse selection and network design under regulated plan prices: Evidence from medicaid. *Journal of Health Economics*, 97:102901.

Liang, A. and Lu, J. (2024a). Algorithmic fairness and social welfare. *arXiv*. arXiv:2404.04424.

Liang, A. and Lu, J. (2024b). Algorithmic fairness and social welfare. *AEA Papers and Proceedings*, 114:628–632.

Mirrlees, J. A. (1971). An exploration in the theory of optimum income taxation. *The Review of Economic Studies*, 38(2):175–208.

Miyagishima, K. (2022). Efficiency, equity, and social rationality under uncertainty. *Economic Theory*, 73:237–255.

Pazner, E. and Schmeidler, D. (1978). Egalitarian equivalent allocations: A new concept of economic equity. *Quarterly Journal of Economics*, 92(4):671–687.

Plaut, B. and Roughgarden, T. (2020). Almost envy-freeness with general valuations. *SIAM Journal on Discrete Mathematics*, pages 401–419.

Rawls, J. (1971). *A Theory of Justice*. Harvard University Press.

Rochet, J.-C. (1987). A necessary and sufficient condition for rationalizability in a quasi-linear context. *Journal of Mathematical Economics*, 16(2):191–200.

Rothschild, C. (2024). Welfare analysis in insurance markets. *The Geneva Risk and Insurance Review*, 49:36–58.

Rothschild, M. and Stiglitz, J. (1976). Equilibrium in competitive insurance markets: An essay on the economics of imperfect information. *The Quarterly Journal of Economics*, 90(4):629–649.

Saez, E. (2001). Using elasticities to derive optimal income tax rates. *The Review of Economic Studies*, 68(1):205–229.

Saez, E. and Stantcheva, S. (2016). Generalized social marginal welfare weights for optimal tax theory. *American Economic Review*, 106(1):24–45.

Saez, E. and Zucman, G. (2019a). Letter regarding senator elizabeth warren's proposed progressive wealth tax. Letter, University of California, Berkeley. January 18, 2019. Available at https://gabriel-zucman.eu/files/saez-zucman-wealthtax-warren.pdf.

Saez, E. and Zucman, G. (2019b). Progressive wealth taxation. *Brookings Papers on Economic Activity*, 50(2):437–533. Fall.

Scheuer, F. and Slemrod, J. (2021). Taxing our wealth. *Journal of Economic Perspectives*, 35(1):207–230.

Shavell, S. and Weiss, L. (1979). The optimal payment of unemployment insurance benefits over time. *Journal of Political Economy*, 87(6):1347–1362.

Sher, I. (2024). Generalized social marginal welfare weights imply inconsistent comparisons of tax policies. *American Economic Review*, 114(11):3551–3577.

Stiglitz, J. E. and Weiss, A. (1981). Credit rationing in markets with imperfect information. *American Economic Review*, 71(3):393–410.

Thomson, W. (2011). Fair allocation rules. In Arrow, K. J., Sen, A., and Suzumura, K., editors, *Handbook of Social Choice and Welfare*, volume 2 of *Handbook of Social Choice and Welfare*, pages 393–506. Elsevier.

Varian, H. R. (1974). Equity, envy, and efficiency. *Journal of Economic Theory*, 9(1):63–91.

Wunsch, C. and Zabrodina, V. (2023). Unemployment insurance with response heterogeneity. CESifo Working Paper 10704, CESifo.

A Technical toolbox for certainty-equivalent gauges

This appendix collects technical facts about certainty-equivalent (CE) gauges that are used repeatedly in the paper. Section A.1 records basic properties of certainty equivalents under expected utility and isolates the cash-additivity condition needed to treat deterministic transfers separately. Section A.2 provides the corresponding properties for a robust (multiple-priors) evaluator. Section A.3 derives continuity bounds for CE-based gain indices and records properties of the insurance value function $V_i(\cdot)$. Section A.4 records compactness and max-gap identities used when optimizing continuous criteria on Pareto frontiers.

Throughout, M is a finite set of states and we identify a lottery with a consumption vector $c = (c_m)_{m\in M} \in \mathbb{R}^M$. For $a \in \mathbb{R}$ we write $a\mathbf{1}$ for the constant vector with all components equal to a.

A.1 Expected-utility certainty equivalents

This subsection fixes an identity $i \in N$ and records the existence and regularity properties of the expected-utility certainty equivalent used in the main text and later appendices.

Assumption 3 (CE admissibility under expected utility)**.** *The Bernoulli utility $v_i : \mathbb{R} \to \mathbb{R}$ is strictly increasing and continuous, and beliefs are given by $p_i \in \Delta(M)$. We restrict attention to $c \in \mathbb{R}^M$ such that*

$$U_i(c) := \sum_{m\in M} p_{i,m} v_i(c_m) \in v_i(\mathbb{R}),$$

so that v_i^{-1} is well-defined at $U_i(c)$. When we use curvature statements, we additionally assume v_i is twice continuously differentiable and strictly concave.

Under Assumption 3, the expected-utility certainty equivalent is

$$\mathrm{CE}_i(c) := v_i^{-1}\big(U_i(c)\big), \qquad \text{equivalently} \qquad v_i(\mathrm{CE}_i(c)) = U_i(c).$$

Lemma 6 (Existence, monotonicity, continuity, and cash additivity)**.** *Maintain Assumption 3. For each admissible $c \in \mathbb{R}^M$:*

1. *$CE_i(c)$ exists and is unique.*
2. *If $c \geq c'$ componentwise, then $CE_i(c) \geq CE_i(c')$.*
3. *The mapping $c \mapsto CE_i(c)$ is continuous on its admissible domain.*

Moreover, $CE_i(a\mathbf{1}) = a$ for all $a \in \mathbb{R}$. If v_i is affine or of CARA form (equivalently, $v_i(x) = A + Be^{-\alpha x}$ with $B \neq 0$ and $\alpha > 0$, up to affine renormalization), then CE_i is cash additive:

$$CE_i(c + a\mathbf{1}) = CE_i(c) + a \quad \text{and} \quad CE_i(a\mathbf{1}) = a \qquad \text{for all } a \in \mathbb{R}. \tag{17}$$

Proof. Existence and uniqueness follow because $U_i(c) \in v_i(\mathbb{R})$ and v_i is strictly increasing and continuous, so v_i^{-1} exists on $v_i(\mathbb{R})$ and $\mathrm{CE}_i(c) = v_i^{-1}(U_i(c))$ is uniquely defined.

If $c \geq c'$ componentwise, then $v_i(c_m) \geq v_i(c'_m)$ for each m, hence $U_i(c) \geq U_i(c')$. Applying the increasing inverse v_i^{-1} yields $\mathrm{CE}_i(c) \geq \mathrm{CE}_i(c')$.

Continuity holds because $c \mapsto U_i(c)$ is continuous (a finite sum of continuous functions) and v_i^{-1} is continuous on $v_i(\mathbb{R})$.

For any $a \in \mathbb{R}$, $U_i(a\mathbf{1}) = \sum_m p_{i,m} v_i(a) = v_i(a)$, so $\mathrm{CE}_i(a\mathbf{1}) = v_i^{-1}(v_i(a)) = a$.

For cash additivity, if v_i is affine then $\mathrm{CE}_i(c) = \sum_m p_{i,m} c_m$ and (17) is immediate. If v_i is CARA (up to affine renormalization), write

$$\mathrm{CE}_i(c) = -\frac{1}{\alpha} \log\left(\sum_{m \in M} p_{i,m} \exp(-\alpha c_m) \right).$$

Then

$$\mathrm{CE}_i(c + a\mathbf{1}) = -\frac{1}{\alpha} \log\left(\sum_m p_{i,m} e^{-\alpha(c_m + a)} \right) = -\frac{1}{\alpha} \log\left(e^{-\alpha a} \sum_m p_{i,m} e^{-\alpha c_m} \right) = \mathrm{CE}_i(c) + a,$$

and the identity $\mathrm{CE}_i(a\mathbf{1}) = a$ is already shown. □

Remark 1 (On deterministic transfers)**.** *When we separate deterministic payments $t_i(g)$ from state-contingent resources, the identities $CE_i(c - t\mathbf{1}) = CE_i(c) - t$ and $\Gamma_i(g) = \Delta_i(g) - t_i(g)$ hold under cash additivity (17). If one prefers to allow general v_i without cash additivity, the analysis can be conducted by absorbing deterministic payments into the consumption vector before applying CE_i, so that the gain index is computed directly from the induced state-contingent payoff.*

A.2 Robust certainty equivalents under ambiguity

This subsection records the multiple-priors extension used when physical uncertainty is evaluated robustly. The only properties used elsewhere are existence of the inner minimizer, monotonicity, continuity, and the same cash-additivity conclusion under affine or CARA utility.

Assumption 4 (Multiple priors)**.** *For each $i \in N$, the set of priors $\Pi_i \subseteq \Delta(M)$ is nonempty and compact.*

Definition 7 (Robust utility and robust certainty equivalent)**.** *Maintain Assumptions 3 and 4. Define*

$$U_i^{\Pi}(c) := \min_{p \in \Pi_i} \sum_{m \in M} p_m \, v_i(c_m),$$

and define the robust certainty equivalent (RCE) by

$$CE_i^{\Pi}(c) := v_i^{-1}\big(U_i^{\Pi}(c)\big),$$

whenever $U_i^{\Pi}(c) \in v_i(\mathbb{R})$.

Lemma 7 (Properties of the robust CE)**.** *Maintain Assumptions 3 and 4. For each* $i \in N$*:*

1. *For every* $c \in \mathbb{R}^M$*, the minimum defining* $U_i^{\Pi}(c)$ *is attained.*
2. *If* $c \geq c'$ *componentwise and both RCEs are well-defined, then* $CE_i^{\Pi}(c) \geq CE_i^{\Pi}(c')$*.*
3. *The mapping* $c \mapsto CE_i^{\Pi}(c)$ *is continuous on its admissible domain.*

Moreover, $CE_i^{\Pi}(a\mathbf{1}) = a$ *for all* $a \in \mathbb{R}$*. If* v_i *is affine or of CARA form (as in Lemma 6), then* CE_i^{Π} *is cash additive:*

$$CE_i^{\Pi}(c + a\mathbf{1}) = CE_i^{\Pi}(c) + a \quad \textit{and} \quad CE_i^{\Pi}(a\mathbf{1}) = a \qquad \textit{for all } a \in \mathbb{R}.$$

Proof. For fixed c, the function $p \mapsto \sum_m p_m v_i(c_m)$ is continuous on $\Delta(M)$ and hence on the compact set Π_i, so it attains a minimum. This proves (1).

If $c \geq c'$ componentwise, then for every $p \in \Pi_i$, $\sum_m p_m v_i(c_m) \geq \sum_m p_m v_i(c'_m)$. Taking minima over $p \in \Pi_i$ preserves the inequality, and applying the increasing inverse v_i^{-1} yields (2) whenever both sides are defined.

The map $(c, p) \mapsto \sum_m p_m v_i(c_m)$ is continuous on $\mathbb{R}^M \times \Pi_i$ and Π_i is compact and independent of c. Berge's maximum theorem implies that $c \mapsto U_i^{\Pi}(c)$ is continuous. Continuity of v_i^{-1} on $v_i(\mathbb{R})$ implies (3) by composition.

For any $a \in \mathbb{R}$, $\sum_m p_m v_i(a) = v_i(a)$ for each $p \in \Pi_i$, hence $U_i^{\Pi}(a\mathbf{1}) = v_i(a)$ and $\mathrm{CE}_i^{\Pi}(a\mathbf{1}) = a$.

For cash additivity, the affine case is immediate. In the CARA case, $v_i(c_m + a)$ introduces a common multiplicative factor inside the minimization exactly as in the expected-utility calculation, and the same logarithmic identity yields the claim. □

A.3 CE gains, net gains, and the insurance value function

This subsection records continuity bounds for CE-based gain indices and the properties of the insurance value function used in Section 7. The arguments are stated for CE_i; the same statements hold with CE_i^{Π} in place of CE_i whenever the RCE is well-defined (Definition 7 and Lemma 7).

Fix a menu G and baseline policy g^0. The CE-based net-gain index used in the paper can be written as

$$\Gamma_i(g) = \mathrm{CE}_i\big(c_i(g)\big) - \mathrm{CE}_i\big(c_i(g^0)\big) - t_i(g),$$

with the convention that deterministic transfers $t_i(g)$ may be absorbed into $c_i(g)$ when preferred (Remark 1).

Lemma 8 (A money-metric Lipschitz bound)**.** *Maintain Assumption 3 and suppose* CE_i *is cash additive (for example,* v_i *is affine or CARA as in Lemma 6). Then for any admissible*

$c, c' \in \mathbb{R}^M$,

$$\left|CE_i(c) - CE_i(c')\right| \le \|c - c'\|_\infty.$$

Consequently, if $g \mapsto c_i(g)$ is continuous and G is compact, then $g \mapsto CE_i(c_i(g))$ is continuous and bounded on G.

Proof. Let $d = \|c - c'\|_\infty$. Then $c' \le c + d\mathbf{1}$ and $c \le c' + d\mathbf{1}$ componentwise. By monotonicity and cash additivity,

$$\mathrm{CE}_i(c') \le \mathrm{CE}_i(c + d\mathbf{1}) = \mathrm{CE}_i(c) + d \quad \text{and} \quad \mathrm{CE}_i(c) \le \mathrm{CE}_i(c' + d\mathbf{1}) = \mathrm{CE}_i(c') + d,$$

which implies $|\mathrm{CE}_i(c) - \mathrm{CE}_i(c')| \le d$. The final statement follows by composition and compactness. □

We now record properties of the insurance value function used in Section 7. Fix income w_i, loss size $\ell > 0$, and loss probability $q_i \in (0, 1)$.

Definition 8 (CE value of coverage)**.** *For identity i, define the CE value of coverage $I \in [0, \ell]$ as*

$$V_i(I) := CE_i\big(w_i,\ w_i - \ell + I\big) - CE_i\big(w_i,\ w_i - \ell\big),$$

where the two-state CE is computed under beliefs $(1 - q_i, q_i)$.

Proposition 5 (Well-definedness and monotonicity of V_i)**.** *Maintain Assumption 3. For each i, the mapping $I \mapsto V_i(I)$ is continuous, satisfies $V_i(0) = 0$, and is strictly increasing on $[0, \ell]$. If, in addition, CE_i is cash additive (e.g. v_i is affine or CARA), then V_i is 1-Lipschitz:*

$$0 \le V_i(I') - V_i(I) \le I' - I \qquad \text{for all } 0 \le I \le I' \le \ell.$$

Proof. Continuity follows from continuity of CE_i (Lemma 6). The identity $V_i(0) = 0$ is immediate.

If $I' > I$, then the loss-state component of $(w_i, w_i - \ell + I')$ strictly exceeds that of $(w_i, w_i - \ell + I)$ while the no-loss component is unchanged. Since v_i is strictly increasing, $U_i(w_i, w_i - \ell + I') > U_i(w_i, w_i - \ell + I)$, hence $\mathrm{CE}_i(w_i, w_i - \ell + I') > \mathrm{CE}_i(w_i, w_i - \ell + I)$ and therefore $V_i(I') > V_i(I)$.

For the Lipschitz bound, let $d := I' - I$ and note that

$$(w_i, w_i - \ell + I') \le (w_i, w_i - \ell + I) + d\mathbf{1}.$$

By monotonicity and cash additivity, $\mathrm{CE}_i(w_i, w_i - \ell + I') \le \mathrm{CE}_i(w_i, w_i - \ell + I) + d$. Subtracting $\mathrm{CE}_i(w_i, w_i - \ell)$ yields $V_i(I') \le V_i(I) + d$. □

The next result gives a derivative formula and a sufficient condition for concavity. The concavity condition is used only for comparative statics and for illustrating curvature of the efficient locus in the insurance application.

Assumption 5 (Smoothness and weak DARA)**.** *For each i, v_i is twice continuously differentiable, strictly increasing and strictly concave, and its Arrow–Pratt absolute risk aversion $A_i(x) := -v_i''(x)/v_i'(x)$ is weakly decreasing in x (weak DARA).*

Proposition 6 (Derivative and concavity of V_i under weak DARA)**.** *Maintain Assumption 5. Let*

$$CE_i(I) := CE_i\big(w_i,\ w_i - \ell + I\big)$$

denote the certainty equivalent of the insured lottery at coverage I. Then V_i is continuously differentiable on $(0, \ell)$ with

$$V_i'(I) = CE_i'(I) = q_i \frac{v_i'(w_i - \ell + I)}{v_i'(CE_i(I))}. \tag{18}$$

If, in addition, CE_i is cash additive (so that Proposition 5 yields the 1-Lipschitz property), then V_i is concave on $[0, \ell]$.

Proof. For $I \in (0, \ell)$, the defining identity

$$v_i(\mathrm{CE}_i(I)) = (1 - q_i)v_i(w_i) + q_i v_i(w_i - \ell + I)$$

and strict monotonicity of v_i imply that $\mathrm{CE}_i(I)$ is well-defined. Differentiating both sides with respect to I yields $v_i'(\mathrm{CE}_i(I))\mathrm{CE}_i'(I) = q_i v_i'(w_i - \ell + I)$, which gives (18). Continuity of V_i' follows from smoothness of v_i.

Differentiate (18) once more. Writing $A_i(x) = -v_i''(x)/v_i'(x)$ and using $v_i'(\mathrm{CE}_i(I)) > 0$, a direct calculation yields

$$\mathrm{CE}_i''(I) = \mathrm{CE}_i'(I)\Big(-A_i(w_i - \ell + I) + \mathrm{CE}_i'(I)\, A_i(\mathrm{CE}_i(I))\Big). \tag{19}$$

Under cash additivity, Proposition 5 implies that $\mathrm{CE}_i(\cdot)$ is 1-Lipschitz on $[0, \ell]$ and therefore $\mathrm{CE}_i'(I) \le 1$ for all $I \in (0, \ell)$.

The insured lottery first-order stochastically dominates the certainty $w_i - \ell + I$ because it equals that value in the loss state and exceeds it in the no-loss state. By monotonicity of CE_i (Lemma 6), this implies $\mathrm{CE}_i(I) \ge w_i - \ell + I$. Under weak DARA, $A_i(\mathrm{CE}_i(I)) \le A_i(w_i - \ell + I)$. Combining this inequality with $\mathrm{CE}_i'(I) \le 1$ implies $\mathrm{CE}_i'(I)A_i(\mathrm{CE}_i(I)) \le A_i(w_i - \ell + I)$, so the bracketed term in (19) is weakly negative and therefore $\mathrm{CE}_i''(I) \le 0$. Thus $\mathrm{CE}_i(\cdot)$ is concave, and so is $V_i(\cdot) = \mathrm{CE}_i(\cdot) - \mathrm{CE}_i(0)$. □

A.4 Compactness, existence, and max-gap identities

This subsection records basic existence and algebraic facts used when optimizing continuous criteria on Pareto frontiers and when working with max-gap objectives in the main text and later appendices.

Lemma 9 (Attainment on compact sets)**.** *Let X be a compact metric space and let $f : X \to \mathbb{R}$ be continuous. Then $\arg\min_{x \in X} f(x)$ and $\arg\max_{x \in X} f(x)$ are nonempty and compact.*

Proof. By continuity, $f(X)$ is compact in $\mathbb{R}$ and therefore attains its minimum and maximum. Closedness of the argmin and argmax sets follows from continuity, and compactness follows because they are closed subsets of a compact set. □

Lemma 10 (Continuity of max-gap functionals)**.** *Let G be a topological space and let $(\Gamma_i)_{i\in N}$ be continuous real-valued functions on G, with N finite. Then the mappings*

$$g \mapsto \max_{i\in N} \Gamma_i(g), \qquad g \mapsto \min_{i\in N} \Gamma_i(g), \qquad g \mapsto D^{CE}(g) := \max_{i,j\in N} \big(\Gamma_j(g) - \Gamma_i(g)\big)$$

are continuous on G.

Proof. For finite N, $g \mapsto \max_i \Gamma_i(g)$ and $g \mapsto \min_i \Gamma_i(g)$ are finite maxima and minima of continuous functions and hence continuous. The map $g \mapsto D^{CE}(g)$ is a finite maximum of continuous functions $g \mapsto \Gamma_j(g) - \Gamma_i(g)$ and is therefore continuous. □

Lemma 11 (Max-gap identities)**.** *For any vector $x = (x_i)_{i\in N} \in \mathbb{R}^N$ with N finite,*

$$\max_{i,j\in N} (x_j - x_i) = \max_{j\in N} x_j - \min_{i\in N} x_i.$$

In particular, if $N = \{L, H\}$, then $\max_{i,j}(x_j - x_i) = |x_L - x_H|$.

Proof. The left-hand side is maximized by choosing j attaining $\max_k x_k$ and i attaining $\min_k x_k$, which yields the right-hand side. The two-group specialization is immediate. □

B Axiomatic characterization of baseline-relative CE gains

This appendix identifies the *baseline-relative* money-metric gain index used in the main text. It shows that once gains are calibrated in dollars by indifference to sure shifts around a baseline, the only admissible gain index is a difference of certainty equivalents. The same argument also clarifies how deterministic payments enter one-for-one when the CE gauge is cash additive.

Fix an identity $i \in N$ and maintain Assumption 1. Outcomes are state-contingent consumption vectors $c \in \mathbb{R}^M$, evaluated relative to a baseline $c^0 \in \mathbb{R}^M$. We use the certainty equivalent $\mathrm{CE}_i(\cdot)$ from Definition 2 and impose cash additivity as a property of the monetary gauge, as in Lemma 6 (for example, under affine or CARA Bernoulli utility). Cash additivity is the restriction that makes sure amounts enter the money metric at face value.

A *baseline-relative gain index* is a mapping $G_i : \mathbb{R}^M \times \mathbb{R}^M \to \mathbb{R}$, where $G_i(c, c^0)$ is interpreted as the sure monetary increment over the baseline that is equivalent, for identity i, to replacing c^0 by c.

Assumption 6 (Baseline-relative money metric and calibration by indifference)**.** *A gain index G_i satisfies:*

1. Baseline normalization: $G_i(c^0, c^0) = 0$ *for all* $c^0 \in \mathbb{R}^M$.

2. Translation invariance (common shift): *for all* $c, c^0 \in \mathbb{R}^M$ *and* $a \in \mathbb{R}$,

$$G_i(c + a\mathbf{1},\, c^0 + a\mathbf{1}) = G_i(c, c^0).$$

3. Calibration by indifference to sure shifts: *for all* $c, c^0 \in \mathbb{R}^M$ *and* $x \in \mathbb{R}$,

$$c \sim_i c^0 + x\mathbf{1} \quad \Longrightarrow \quad G_i(c, c^0) = x,$$

where $\sim_i$ *denotes indifference under the preference in Assumption 1.*

Assumption 6(2) says that a common sure shift of both the baseline and the reform does not change the gain. Assumption 6(3) fixes the money scale by requiring the index to agree with sure amounts whenever the reform is equivalent to a sure shift of the baseline.

Proposition 7 (Uniqueness of CE gains as a baseline-relative money metric)**.** *Maintain Assumption 1 and suppose* CE_i *is cash additive as in Lemma 6. Let* G_i *be any gain index satisfying Assumption 6. Then, for all* $c, c^0 \in \mathbb{R}^M$,

$$G_i(c, c^0) \;=\; CE_i(c) - CE_i(c^0).$$

Proof. Fix (c, c^0) and set $x := \mathrm{CE}_i(c) - \mathrm{CE}_i(c^0)$. Cash additivity implies

$$\mathrm{CE}_i(c^0 + x\mathbf{1}) = \mathrm{CE}_i(c^0) + x = \mathrm{CE}_i(c).$$

By definition of the certainty equivalent, $\mathrm{CE}_i(c) = \mathrm{CE}_i(c^0 + x\mathbf{1})$ implies $c \sim_i c^0 + x\mathbf{1}$. Assumption 6(3) then yields $G_i(c, c^0) = x$. □

Proposition 7 is the baseline-relative counterpart of Proposition 1 in the main text. Proposition 1 pins down the certainty-equivalent *level* index as the unique preference-consistent money metric for risky prospects. Proposition 7 pins down the unique *difference* representation of gains relative to an arbitrary baseline once the gain index is calibrated in dollars by indifference to sure shifts.

Lemma 12 (Deterministic transfers enter additively)**.** *Maintain Assumption 1 and suppose* CE_i *is cash additive. For any* $c \in \mathbb{R}^M$ *and deterministic* $t \in \mathbb{R}$,

$$CE_i(c - t\mathbf{1}) = CE_i(c) - t.$$

Consequently, if $c_i(g) = \tilde{c}_i(g) - t_i(g)\mathbf{1}$ *and* $c_i(g^0) = \tilde{c}_i(g^0) - t_i(g^0)\mathbf{1}$*, then*

$$CE_i(c_i(g)) - CE_i(c_i(g^0)) = \Big(CE_i(\tilde{c}_i(g)) - CE_i(\tilde{c}_i(g^0))\Big) - \big(t_i(g) - t_i(g^0)\big).$$

Proof. Cash additivity gives $\mathrm{CE}_i(c - t\mathbf{1}) = \mathrm{CE}_i(c) + (-t) = \mathrm{CE}_i(c) - t$. □

Remark 2 (Robust CE version)**.** *All statements in this appendix extend verbatim if one replaces* CE_i *by the robust certainty equivalent* CE_i^{Π} *from Definition 7, provided indifference* $\sim_i$ *in Assumption 6(3) is interpreted under the robust preference induced by* U_i^{Π}*. The*

proofs use only calibration by indifference to sure shifts and cash additivity of the certainty-equivalent gauge.

C CE units and the surrogate selector: bridge results

This appendix supplies the technical bridge used in Sections 5–6. The companion paper constructs a complaint-based selector on the Pareto frontier from within-menu gauges through *directed slacks*. The present paper specializes those gauges to baseline-relative certainty-equivalent (CE) net-gain dollars. Under this specialization, the selector admits an explicit and economically transparent representation as a max-gap minimization program on $\mathrm{PE}(G)$ with a leximin refinement, and with two identities it reduces to minimizing a single absolute CE net-gain gap.

Throughout, fix a menu G and a baseline policy g^0. For each identity $i \in N$, let $\Gamma_i(g)$ denote the CE net gain from g relative to g^0 (Definition 3). For clarity, deterministic transfers are treated as already absorbed into the state-contingent payoff in this appendix. When one wishes to separate transfers explicitly, the bridge remains valid under the cash-additivity condition discussed in Lemma 12 (Appendix B) and Remark 1 (Appendix A).

Proposition 8 (Epigraph and dual representations of the CE continuation program)**.** *Fix a menu G and baseline g^0. Let $\Gamma_i(g)$ denote the baseline-relative CE net gain for identity $i \in N$ (Definition 3), and let*

$$D^{CE}(g) = \max_{i,j\in N}\big(\Gamma_j(g) - \Gamma_i(g)\big) = \max_{j\in N}\Gamma_j(g) - \min_{i\in N}\Gamma_i(g)$$

be the maximal CE gap (Definition 5).

(i) For any fixed policy $g \in G$,

$$D^{CE}(g) = \min_{\alpha\in\mathbb{R}}\ \min_{t\geq 0}\Big\{t:\ \alpha \leq \Gamma_i(g) \leq \alpha + t \text{ for all } i \in N\Big\}.$$

Moreover, at any minimizer (α, t) one has $t = D^{CE}(g)$ and $\alpha \in \big[\min_i \Gamma_i(g),\ \max_i \Gamma_i(g)\big]$.

(ii) Consequently, the CE-indexed continuation problem can be written as the epigraph program

$$\min_{g\in PE(G),\ \alpha\in\mathbb{R},\ t\geq 0}\ t \quad s.t. \quad \alpha \leq \Gamma_i(g) \leq \alpha + t \ \text{ for all } i \in N.$$

At any solution $(g^\star, \alpha^\star, t^\star)$, the slack level satisfies $t^\star = D^{CE}(g^\star)$.

(iii) For any fixed $g \in G$, the maximal CE gap admits the dual representation

$$D^{CE}(g) = \max_{u,v\in\Delta(N)}\ \sum_{i\in N}\big(u_i - v_i\big)\Gamma_i(g),$$

where $\Delta(N)$ is the simplex on N. The maximum is attained by choosing u supported

on $\arg\max_i \Gamma_i(g)$ *and* v *supported on* $\arg\min_i \Gamma_i(g)$. *Equivalently,*

$$\min_{g\in PE(G)} D^{CE}(g) = \min_{g\in PE(G)} \max_{u,v\in\Delta(N)} \sum_{i\in N} \big(u_i - v_i\big)\Gamma_i(g).$$

Proof. For (i), fix g and write $m(g) := \min_i \Gamma_i(g)$ and $M(g) := \max_i \Gamma_i(g)$. If (α, t) is feasible, then $\alpha \le m(g)$ and $M(g) \le \alpha + t$, hence $t \ge M(g) - \alpha \ge M(g) - m(g) = D^{CE}(g)$. Conversely, taking $\alpha = m(g)$ and $t = M(g) - m(g)$ is feasible, so the minimum equals $D^{CE}(g)$. The stated bounds on α follow from feasibility.

For (ii), apply (i) pointwise and minimize over $g \in \mathrm{PE}\,(G)$.

For (iii), use the standard simplex identities

$$\max_{u\in\Delta(N)} \sum_i u_i\Gamma_i(g) = \max_i \Gamma_i(g) = M(g), \qquad \max_{v\in\Delta(N)} \sum_i v_i\big(-\Gamma_i(g)\big) = -\min_i \Gamma_i(g) = -m(g),$$

so $D^{CE}(g) = M(g) - m(g) = \max_{u,v\in\Delta(N)} \sum_i (u_i - v_i)\Gamma_i(g)$. Attainment holds by taking u and v as unit masses on maximizers and minimizers, respectively. The final display is immediate by minimizing over $g \in \mathrm{PE}\,(G)$. □

C.1 Directed slacks and the directed-slack leximin rule (import from the companion)

The companion paper studies a menu of sure outcomes in $\mathbb{R}^N$ together with a profile of strictly increasing gauges $\tilde{u} = (\tilde{u}_i)_{i\in N}$, where $\tilde{u}_i : \mathbb{R} \to \mathbb{R}$ is applied to coordinates. For an outcome vector $x = (x_k)_{k\in N} \in \mathbb{R}^N$, the directed slack from i to j is

$$d_{ij}(x) \;:=\; \tilde{u}_i(x_i) - \tilde{u}_i(x_j), \qquad i \neq j. \tag{20}$$

Let $\mathbf{d}(x)$ collect $\{d_{ij}(x) : i \neq j\}$ and let $\mathbf{d}^{\uparrow}(x)$ denote the same vector sorted in weakly increasing order, so the first entry is the worst directed slack. The directed-slack leximin correspondence selects Pareto-undominated outcomes whose sorted slack vector is lexicographically maximal:

$$\mathrm{LexMin}^{\,\mathrm{dir}}(G;\tilde{u}) := \Big\{x \in \mathrm{PE}\,(G) :\; \mathbf{d}^{\uparrow}(x) \text{ is lexicographically maximal over } \mathrm{PE}\,(G)\Big\}. \tag{21}$$

In this appendix we use only that $\mathrm{LexMin}^{\,\mathrm{dir}}(G;\tilde{u})$ selects Pareto-efficient outcomes and that it implements the companion paper's continuation logic by maximizing the worst directed slack and then refining by the next-worst slacks in the leximin order.

C.2 From policies to a sure-outcome menu in CE dollars

Define the CE image of the policy menu by

$$\Gamma(G) \;:=\; \big\{\Gamma(g) := (\Gamma_i(g))_{i\in N} \in \mathbb{R}^N :\; g \in G\big\}.$$

Because $\Gamma_i(g)$ is constructed from identity i's preferences and the baseline is fixed across g, Pareto comparisons can be made equivalently in the original policy menu and in its CE image.

Lemma 13 (Pareto efficiency is preserved under the CE map)**.** *Maintain Assumption 1 and absorb deterministic transfers into consumption. For $g, g' \in G$ and each $i \in N$,*

$$c_i(g) \succeq_i c_i(g') \quad \Longleftrightarrow \quad \Gamma_i(g) \geq \Gamma_i(g').$$

Consequently,

$$g \in PE(G) \quad \Longleftrightarrow \quad \Gamma(g) \text{ is Pareto undominated in } \Gamma(G) \text{ under componentwise order.}$$

Proof. Under Assumption 1, $c_i(g) \succeq_i c_i(g')$ if and only if $U_i(c_i(g)) \geq U_i(c_i(g'))$. Since v_i is strictly increasing, this is equivalent to $\mathrm{CE}_i(c_i(g)) \geq \mathrm{CE}_i(c_i(g'))$. Subtracting the constant $\mathrm{CE}_i(c_i(g^0))$ yields $\Gamma_i(g) \geq \Gamma_i(g')$. The statement on Pareto undominance follows immediately. □

C.3 CE-directed slacks and explicit max-gap programs

Once outcomes are expressed in CE net-gain dollars, the present paper uses the identity gauge on those dollars,

$$\tilde{u}_i(x) = x \qquad \text{for all } i \in N. \tag{22}$$

Under (22), directed slacks become simple CE-dollar differences.

Definition 9 (Directed CE slack)**.** *For $g \in G$ and $i \neq j$, define the directed CE slack by*

$$d_{ij}^{CE}(g) \;:=\; \Gamma_i(g) - \Gamma_j(g).$$

Let $\mathbf{d}^{CE}(g)$ collect all directed CE slacks and let $\mathbf{d}^{CE,\uparrow}(g)$ denote the sorted vector (worst first).

The continuation objective can be expressed directly in terms of the range of CE net gains.

Lemma 14 (Worst directed CE slack equals minus the maximal CE gap)**.** *For any $g \in G$,*

$$\min_{i \neq j} d_{ij}^{CE}(g) \;=\; -\max_{i,j \in N} \big(\Gamma_j(g) - \Gamma_i(g)\big) \;=\; -D^{CE}(g),$$

where $D^{CE}(g)$ is the maximal CE gap from Definition 5.

Proof. Apply Lemma 11 to the vector $\Gamma(g) \in \mathbb{R}^N$. □

The next proposition records the explicit form of the continuation program and the role of the leximin refinement. The assumptions ensure that the relevant minima and maxima are attained on $\mathrm{PE}(G)$.

Proposition 9 (CE-indexed directed-slack leximin and the explicit CE continuation)**.** *Assume $PE(G)$ is nonempty and compact and that $g \mapsto \Gamma(g)$ is continuous on $PE(G)$.*

1. *The directed-slack leximin correspondence on the CE image selects exactly the CE images of Pareto-efficient policies that minimize the maximal CE gap:*

$$LexMin^{\mathrm{dir}}\big(\Gamma(G);\mathrm{id}\big) = \Big\{\Gamma(g) :\ g \in \arg\min_{h\in PE(G)} D^{CE}(h)\Big\}.$$

2. *The maximization of the sorted slack vector refines the max-gap minimization as follows. It first minimizes $D^{CE}(g)$ on $PE(G)$, and among minimizers it applies the leximin order to the full directed-slack vector $\mathbf{d}^{CE}(g)$.*

3. *If $PEEF^{CE}(G) \neq \varnothing$, then $D^{CE}(g) = 0$ at every $g \in PEEF^{CE}(G)$ and every selected policy is CE envy-free and Pareto efficient.*

Proof. Existence of minimizers of D^{CE} on $\mathrm{PE}(G)$ follows from Lemma 10 and Lemma 9. Lemma 14 implies that minimizing D^{CE} is equivalent to maximizing the worst directed CE slack $\min_{i\neq j} d_{ij}^{CE}(g)$. The correspondence (21) compares the entire sorted directed-slack vector and therefore refines this maximin objective by successively maximizing the second-worst slack, then the third-worst, and so on, which yields the set equality in (i) and the refinement description in (ii). For (iii), $D^{CE}(g) = 0$ holds if and only if all CE net gains are equal, and this is equivalent to CE envy-freeness by Lemma 2. □

C.4 Two-identity reduction

When there are two identities, the maximal CE gap collapses to a single absolute difference. This is the simplification used in Corollary 1 and Corollary 2.

Lemma 15 (Two identities: minimizing range equals minimizing the absolute gap)**.** *If $N = \{1, 2\}$, then for any $g \in G$,*

$$D^{CE}(g) = |\Gamma_1(g) - \Gamma_2(g)|.$$

Consequently,

$$\arg\min_{g\in PE(G)} D^{CE}(g) = \arg\min_{g\in PE(G)} |\Gamma_1(g) - \Gamma_2(g)|.$$

Proof. This is the two-component specialization of Lemma 11. □

The conclusion of this appendix can be summarized succinctly. Expressing a policy menu in baseline-relative CE net-gain dollars turns the companion paper's directed-slack continuation into the explicit program of minimizing the maximal CE net-gain disparity on $\mathrm{PE}(G)$, with the same leximin refinement used throughout the paper.

D Veil-of-ignorance foundations and ambiguity over identity

This appendix provides proofs and auxiliary results for Section 6. The first subsection shows that worst-case veil regret depends only on the range of certainty-equivalent (CE) net gains. The second subsection records a compact verification of the veil equivalence result under the existence and continuity conditions used elsewhere. The third subsection gives a complementary foundation in which the evaluator is ambiguity-averse about the identity distribution. The last subsection records a leximin refinement stated directly in terms of regret profiles and relates it to the refinement used in the companion selector.

Throughout, fix a menu G and baseline g^0. For each identity $i \in N$, let $\Gamma_i(g)$ denote the CE net gain induced by g relative to g^0 (Definition 3). Write $\Gamma(g) := (\Gamma_i(g))_{i\in N} \in \mathbb{R}^N$.

D.1 Worst-case veil regret equals the maximal CE gap

Recall from Definition 6 that for $g \in G$ and $i \in N$,

$$R_i(g) := \max_{j\in N}\big(\Gamma_j(g) - \Gamma_i(g)\big), \qquad R(g) := \max_{i\in N} R_i(g).$$

Lemma 16 (Regret representation)**.** *For every $g \in G$,*

$$R_i(g) = \max_{j\in N} \Gamma_j(g) - \Gamma_i(g) \quad \textit{for each } i \in N, \qquad R(g) = D^{CE}(g) = \max_{j\in N} \Gamma_j(g) - \min_{i\in N} \Gamma_i(g),$$

where $D^{CE}(g)$ is the maximal CE gap from Definition 5. In particular, $R(g) = 0$ if and only if g is CE envy-free.

Proof. Fix g. For each i, the term $\Gamma_i(g)$ does not depend on the maximization index j, so

$$R_i(g) = \max_{j\in N}\big(\Gamma_j(g) - \Gamma_i(g)\big) = \max_{j\in N} \Gamma_j(g) - \Gamma_i(g).$$

Taking the maximum over i yields

$$R(g) = \max_{i\in N}\Big(\max_{j\in N} \Gamma_j(g) - \Gamma_i(g)\Big) = \max_{j\in N} \Gamma_j(g) - \min_{i\in N} \Gamma_i(g).$$

By Lemma 11 (Appendix A), the right-hand side equals $\max_{i,j}(\Gamma_j(g) - \Gamma_i(g)) = D^{CE}(g)$. Finally, $R(g) = 0$ holds if and only if $\max_j \Gamma_j(g) = \min_i \Gamma_i(g)$, which is equivalent to $\Gamma_i(g)$ being constant across i, hence to CE envy-freeness by Lemma 2. □

D.2 Proof of the veil equivalence theorem

This subsection verifies Theorem 2 under the regularity conditions ensuring that minimizers exist on the Pareto frontier.

Theorem 5 (Veil equivalence (detailed version of Theorem 2))**.** *Assume $PE(G)$ is nonempty and compact and that $g \mapsto \Gamma(g)$ is continuous on $PE(G)$. Then:*

1. *If $PEEF^{CE}(G) \neq \varnothing$, then*

$$V^{CE}(G) = PEEF^{CE}(G),$$

and under any single-valued refinement on $PEEF^{CE}(G)$ one has $V^{CE}(G) = F^{CE}(G) = \{g^{CE}\}$, where g^{CE} is the CE equalizer from Proposition 3.

2. *If $PEEF^{CE}(G) = \varnothing$, then*

$$V^{CE}(G) = \arg\min_{g \in PE(G)} D^{CE}(g) = \arg\max_{g \in PE(G)} S^{CE}(g),$$

and with the shared refinement rule one has $V^{CE}(G) = F^{CE}(G)$.

Proof. By Lemma 16, $R(g) = D^{CE}(g)$ for all $g \in \mathrm{PE}(G)$. By Lemma 10 (Appendix A), the map $g \mapsto D^{CE}(g)$ is continuous on $\mathrm{PE}(G)$. Since $\mathrm{PE}(G)$ is compact, Lemma 9 (Appendix A) implies that $\arg\min_{\mathrm{PE}(G)} D^{CE}(g)$ is nonempty, so the veil correspondence $V^{CE}(G) = \arg\min_{g \in \mathrm{PE}(G)} R(g)$ is well-defined and equals $\arg\min_{\mathrm{PE}(G)} D^{CE}(g)$.

If $\mathrm{PEEF}^{CE}(G) \neq \varnothing$, then $D^{CE}(g) = 0$ for every $g \in \mathrm{PEEF}^{CE}(G)$ by Lemma 16, and $D^{CE}(g) \geq 0$ for all g by definition. Hence the minimum value of D^{CE} on $\mathrm{PE}(G)$ is 0 and the set of minimizers is exactly $\mathrm{PEEF}^{CE}(G)$, proving the first claim. Under any single-valued refinement on $\mathrm{PEEF}^{CE}(G)$, Proposition 3 delivers the unique CE equalizer g^{CE}, hence $V^{CE}(G) = F^{CE}(G) = \{g^{CE}\}$.

If $\mathrm{PEEF}^{CE}(G) = \varnothing$, then no $g \in \mathrm{PE}(G)$ attains $D^{CE}(g) = 0$, so the minimax value is strictly positive and

$$V^{CE}(G) = \arg\min_{g \in \mathrm{PE}(G)} D^{CE}(g) = \arg\max_{g \in \mathrm{PE}(G)} S^{CE}(g),$$

where $S^{CE}(g) = -D^{CE}(g)$. Theorem 1 identifies $F^{CE}(G)$ with this continuation problem on $\mathrm{PE}(G)$, and the shared refinement delivers $V^{CE}(G) = F^{CE}(G)$. □

D.3 Ambiguity over identity distributions

This subsection records a complementary foundation in which the evaluator is ambiguity-averse about the distribution over identities. The only object needed is the CE net-gain profile $\Gamma(g)$, so the analysis is independent of how CE net gains are computed from physical uncertainty.

Definition 10 (Identity lotteries and robust veil value)**.** *Given a distribution $\mu \in \Delta(N)$ over identities, define the μ-veil value of policy g by*

$$W_\mu(g) := \sum_{i \in N} \mu_i \, \Gamma_i(g).$$

Given a nonempty compact set of priors $M \subseteq \Delta(N)$, define the robust veil value by

$$W_M(g) := \min_{\mu \in M} W_\mu(g) = \min_{\mu \in M} \sum_{i \in N} \mu_i \, \Gamma_i(g).$$

Proposition 10 (Robust identity uncertainty yields a maximin CE criterion)**.** *Suppose $M \subseteq \Delta(N)$ is nonempty and contains every point-mass distribution δ^i for $i \in N$. Then for every policy g,*

$$W_M(g) = \min_{i \in N} \Gamma_i(g).$$

Consequently, maximizing $W_M(g)$ over a feasible set selects the policy that maximizes the minimum CE net gain across identities.

Proof. For any $\mu \in \Delta(N)$, the value $W_\mu(g) = \sum_i \mu_i \Gamma_i(g)$ is a convex combination of the components of $\Gamma(g)$, hence $W_\mu(g) \geq \min_i \Gamma_i(g)$. Minimizing over $\mu \in \mathrm{M}$ yields $W_{\mathrm{M}}(g) \geq \min_i \Gamma_i(g)$. If $i^\star \in \arg\min_i \Gamma_i(g)$, then $\delta^{i^\star} \in \mathrm{M}$ implies $W_{\mathrm{M}}(g) \leq W_{\delta^{i^\star}}(g) = \Gamma_{i^\star}(g) = \min_i \Gamma_i(g)$, which proves the equality. □

D.4 Leximin refinements and tie-breaking

When the minimax regret correspondence $V^{CE}(G)$ is not single-valued, it is natural to refine by comparing the entire regret profile. In this subsection, write $\mathbf{R}(g) := (R_i(g))_{i \in N} \in \mathbb{R}^N_+$ for the regret vector, and write $\mathbf{R}^{\downarrow}(g)$ for the same vector sorted in weakly decreasing order.

Definition 11 (Veil leximin refinement)**.** *Define the veil leximin correspondence by*

$$V^{CE,lex}(G) := \Big\{ g \in PE(G) : \ \mathbf{R}^{\downarrow}(g) \text{ is lexicographically minimal over } PE(G) \Big\}.$$

Proposition 11 (Leximin refines minimax regret)**.** *Every $g \in V^{CE,lex}(G)$ minimizes $R(g)$ on $PE(G)$, hence $V^{CE,lex}(G) \subseteq V^{CE}(G)$. If $N = \{1, 2\}$, then $V^{CE,lex}(G) = V^{CE}(G)$.*

Proof. If $g \in V^{CE,\mathrm{lex}}(G)$, then $\mathbf{R}^{\downarrow}(g)$ is lexicographically minimal on $\mathrm{PE}(G)$, so its first component is minimal on $\mathrm{PE}(G)$. The first component of $\mathbf{R}^{\downarrow}(g)$ is $\max_{i \in N} R_i(g) = R(g)$ by definition of sorting. Hence g minimizes $R(g)$ on $\mathrm{PE}(G)$, which implies $V^{CE,\mathrm{lex}}(G) \subseteq V^{CE}(G)$.

If $N = \{1, 2\}$, then $(R_1(g), R_2(g)) = \big((\Gamma_2(g) - \Gamma_1(g))_+, (\Gamma_1(g) - \Gamma_2(g))_+\big)$, so $\mathbf{R}^{\downarrow}(g) = (|\Gamma_1(g) - \Gamma_2(g)|, 0)$ by Lemma 11 and Lemma 16. Therefore, on $\mathrm{PE}(G)$ lexicographic minimization of $\mathbf{R}^{\downarrow}(g)$ is equivalent to minimizing $|\Gamma_1(g) - \Gamma_2(g)|$, which is the minimax criterion. □

Remark 3 (Relation to the refinement used in the paper)**.** *The main text uses the refinement inherited from the companion selector, which applies a leximin order to the directed-slack profile. Appendix C shows that under CE net-gain units the directed-slack profile is the collection of pairwise differences in $\Gamma(g)$. Definition 11 provides a veil-based refinement stated in terms of regret profiles. Both refinements agree with minimax regret, and Proposi-*

tion 11 shows that they coincide in the two-identity environments that drive the insurance application.

E Insurance: proofs, frontier characterizations, and comparative statics

This appendix proves the results used in the insurance application (Section 7). It first verifies the properties of the CE value-of-coverage function $V_i(\cdot)$ and records CARA comparative statics used for interpretation. It then characterizes the one-dimensional efficient locus induced by cap-financed mandates and establishes the monotone incidence property that drives the two-group reduction. It finally studies pooling feasibility under voluntary enrollment and records sufficient conditions and comparative statics that are useful for applications. Generic certainty-equivalent facts are taken from Appendix A.

E.1 Standing notation

There are two groups $i \in \{L, H\}$ with primitives satisfying (3). A policy $g = (I, t_L, t_H)$ specifies coverage $I \in [0, \ell]$ and deterministic payments t_i. Under g,

$$c_i(g) = \big(w_i - t_i,\ w_i - \ell + I - t_i\big), \qquad c_i^0 = (w_i,\ w_i - \ell),$$

and CE net gains take the form

$$\Gamma_i(g) = V_i(I) - t_i, \qquad V_i(I) := \mathrm{CE}_i(w_i,\ w_i - \ell + I) - \mathrm{CE}_i(w_i,\ w_i - \ell).$$

When we use closed forms and derivatives, we impose CARA utility (4), in which case $V_i(\cdot)$ admits (7).

E.2 Properties of $V_i(I)$

Proof of Lemma 4. Continuity and strict monotonicity of $I \mapsto V_i(I)$ follow from Proposition 5 in Appendix A. When curvature is used, the concavity claim follows from Proposition 6. Under CARA, the closed form (7) is given in Example 2. □

The comparative statics stated in the main text under CARA are most transparent at full coverage, where $V_i(\ell)$ reduces to a log-moment expression.

Lemma 17 (CARA derivatives of $V_i(\ell)$)**.** *Maintain CARA (4). Let*

$$B_i := (1 - q_i) + q_i e^{a_i \ell}.$$

Then $V_i(\ell) = \frac{1}{a_i}\ln(B_i)$ *and*

$$\frac{\partial V_i(\ell)}{\partial q_i} = \frac{e^{a_i\ell} - 1}{a_i\, B_i} > 0, \qquad \frac{\partial V_i(\ell)}{\partial a_i} = -\frac{1}{a_i^2}\ln(B_i) + \frac{q_i \ell e^{a_i\ell}}{a_i B_i}.$$

Moreover, $\partial V_i(\ell)/\partial a_i > 0$ *for all* $q_i \in (0,1)$ *and* $\ell > 0$.

Proof. The displayed derivatives follow by direct differentiation of $V_i(\ell) = \frac{1}{a_i}\ln(B_i)$.

For the sign of $\partial V_i(\ell)/\partial a_i$, let Z be the binary random variable equal to ℓ with probability q_i and 0 otherwise. Then $B_i = \mathbb{E}[e^{a_i Z}]$ and $V_i(\ell) = k(a_i)/a_i$ where $k(a) := \ln \mathbb{E}[e^{aZ}]$. Since Z is nondegenerate and bounded, k is strictly convex on $(0,\infty)$ and satisfies $k(0) = 0$. Strict convexity and $k(0) = 0$ imply that $a \mapsto k(a)/a$ is strictly increasing on $(0,\infty)$, which is equivalent to $\partial V_i(\ell)/\partial a_i > 0$. □

E.3 Mandated public insurance with restricted incidence

In the mandated analysis of Section 7, coverage is fixed at $I = \ell$ and incidence is indexed by the contribution ceiling $\bar{w} \in [w_L, w_H]$ through (8)–(10). Along this family,

$$\Gamma_i(\bar{w}) := \Gamma_i\big(g^{\text{man}}(\bar{w})\big) = V_i(\ell) - t_i(\bar{w}), \tag{23}$$

with $t_i(\bar{w})$ given by (10).

Proposition 12 (Frontier in CE net-gain space under the mandate)**.** *Consider the mandated menu* $G^{\text{man}} = \{g^{\text{man}}(\bar{w}) : \bar{w} \in [w_L, w_H]\}$. *Then* $PE(G^{\text{man}}) = G^{\text{man}}$ *and the image frontier in CE net gains is the one-dimensional curve* $\bar{w} \mapsto (\Gamma_L(\bar{w}), \Gamma_H(\bar{w}))$ *defined by* (23) *and* (10). *Moreover,* $\Gamma_L(\bar{w})$ *is strictly increasing and* $\Gamma_H(\bar{w})$ *is strictly decreasing in* $\bar{w}$.

Proof. Fix $\bar{w} < \bar{w}'$. By (10), $t_H(\bar{w}') > t_H(\bar{w})$ and $t_L(\bar{w}') < t_L(\bar{w})$. Since $V_i(\ell)$ is independent of $\bar{w}$, it follows that $\Gamma_H(\bar{w}') < \Gamma_H(\bar{w})$ and $\Gamma_L(\bar{w}') > \Gamma_L(\bar{w})$. Hence neither policy Pareto dominates the other inside G^{man}, so every policy in the mandated menu is Pareto efficient. The stated monotonicity of Γ_L and Γ_H follows from the same inequalities. □

Proof of Lemma 5. Along G^{man},

$$\Gamma_L(\bar{w}) - \Gamma_H(\bar{w}) = \big(V_L(\ell) - V_H(\ell)\big) - \big(t_L(\bar{w}) - t_H(\bar{w})\big).$$

The term $V_L(\ell) - V_H(\ell)$ is constant in $\bar{w}$. By (10), the difference $t_H(\bar{w}) - t_L(\bar{w})$ is strictly increasing in $\bar{w}$, so $t_L(\bar{w}) - t_H(\bar{w})$ is strictly decreasing. Therefore $\Gamma_L(\bar{w}) - \Gamma_H(\bar{w})$ is strictly increasing in $\bar{w}$. □

Equalization in CE net gains requires $\Gamma_L(\bar{w}) = \Gamma_H(\bar{w})$, equivalently

$$V_L(\ell) - V_H(\ell) = t_L(\bar{w}) - t_H(\bar{w}), \tag{24}$$

which reduces, using (10), to

$$V_L(\ell) - V_H(\ell) = (q_L + q_H)\ell \cdot \frac{w_L - \bar{w}}{w_L + \bar{w}}. \tag{25}$$

Lemma 18 (Unique equalizing ceiling (when it exists))**.** *Let* $A := V_L(\ell) - V_H(\ell)$ *and* $B := (q_L + q_H)\ell > 0$. *If there exists* $\bar{w}^* \in [w_L, w_H]$ *satisfying* (25)*, then it is unique. Moreover, any interior solution satisfies*

$$\bar{w}^* = w_L \cdot \frac{B - A}{B + A}, \tag{26}$$

whenever $B + A \neq 0$.

Proof. Uniqueness follows from Lemma 5, since a strictly monotone function can cross zero at most once. Solving (25) for $\bar{w}$ yields (26) whenever $B + A \neq 0$. □

Proof of Theorem 3. Proposition 12 shows that the mandated menu is already Pareto efficient and is parameterized by the scalar $\bar{w} \in [w_L, w_H]$. Lemma 5 shows that the CE net-gain difference $\Gamma_L(\bar{w}) - \Gamma_H(\bar{w})$ is strictly increasing in $\bar{w}$. Corollary 1 therefore applies to the parameterization $\bar{w} \mapsto g^{\text{man}}(\bar{w})$.

If there exists $\bar{w}^*$ with $\Gamma_L(\bar{w}^*) = \Gamma_H(\bar{w}^*)$, strict monotonicity implies uniqueness and the CE-indexed rule selects $g^{\text{man}}(\bar{w}^*)$. If no such $\bar{w}^*$ exists, strict monotonicity implies that the absolute gap $|\Gamma_L(\bar{w}) - \Gamma_H(\bar{w})|$ is minimized at a boundary point, so the CE-indexed rule selects the boundary ceiling stated in Theorem 3. Lemma 18 provides the interior closed form when the equalizer exists. □

When the equalizing ceiling exists and is interior, it is pinned down by (25), equivalently by (26). The next proposition records comparative statics in terms of the primitives that shift CE valuations and expected costs.

Proposition 13 (Comparative statics of the equalizing ceiling)**.** *Suppose the equalizing ceiling* $\bar{w}^* \in (w_L, w_H)$ *exists and is given by* (26)*, with* $A := V_L(\ell) - V_H(\ell)$ *and* $B := (q_L + q_H)\ell$ *satisfying* $B + A \neq 0$. *Then:*

1. $\dfrac{\partial \bar{w}^*}{\partial A} = -\dfrac{2Bw_L}{(B+A)^2} < 0.$
2. $\dfrac{\partial \bar{w}^*}{\partial B} = \dfrac{2Aw_L}{(B+A)^2}$, *so* $\partial \bar{w}^*/\partial B$ *has the sign of* A.
3. *Under CARA,* $\partial A/\partial q_L > 0$ *and* $\partial A/\partial q_H < 0$, *so (holding other primitives fixed)* $\bar{w}^*$ *decreases in* q_L *and increases in* q_H.
4. *Under CARA,* $\partial A/\partial a_L > 0$ *and* $\partial A/\partial a_H < 0$, *so* $\bar{w}^*$ *decreases in* a_L *and increases in* a_H.

Proof. Differentiate (26) with respect to A and B to obtain the derivative expressions in parts (1) and (2). Parts (3) and (4) follow from Lemma 17, since $A = V_L(\ell) - V_H(\ell)$. □

E.4 Voluntary enrollment: pooling feasibility and incidence trade-offs

Under voluntary enrollment, pooled policies satisfy budget balance (12) and participation constraints (13). Under (12) and $\bar{w} \in [w_L, w_H]$, budget balance implies

$$t_L(I, \bar{w}) = \frac{(q_L + q_H)I\, w_L}{w_L + \bar{w}}, \qquad t_H(I, \bar{w}) = \frac{(q_L + q_H)I\, \bar{w}}{w_L + \bar{w}}. \tag{27}$$

Participation constraints (13) become

$$V_L(I) \geq t_L(I, \bar{w}), \qquad V_H(I) \geq t_H(I, \bar{w}). \tag{28}$$

Lemma 19 (Pooling feasibility as an interval of ceilings)**.** *Fix coverage $I \in (0, \ell]$ and suppose $V_i(I) > 0$ for $i \in \{L, H\}$. Define*

$$\underline{\bar{w}}(I) := w_L\Big(\frac{(q_L + q_H)I}{V_L(I)} - 1\Big)_+, \qquad \overline{\bar{w}}(I) := \begin{cases} \dfrac{w_L V_H(I)}{(q_L + q_H)I - V_H(I)} & \text{if } V_H(I) < (q_L + q_H)I, \\ +\infty & \text{if } V_H(I) \geq (q_L + q_H)I. \end{cases}$$

Then (28) holds if and only if

$$\bar{w} \in [\underline{\bar{w}}(I),\, \overline{\bar{w}}(I)] \cap [w_L, w_H]. \tag{29}$$

In particular, for fixed I, pooling with coverage I is feasible if and only if the intersection in (29) is nonempty.

Proof. Using (27), the constraint $V_L(I) \geq t_L(I, \bar{w})$ is equivalent to

$$V_L(I) \geq (q_L + q_H)I\, \frac{w_L}{w_L + \bar{w}} \quad \Longleftrightarrow \quad \bar{w} \geq w_L\Big(\frac{(q_L + q_H)I}{V_L(I)} - 1\Big),$$

and intersecting with $\bar{w} \geq 0$ yields $\underline{\bar{w}}(I)$.

Similarly, $V_H(I) \geq t_H(I, \bar{w})$ is

$$V_H(I) \geq (q_L + q_H)I\, \frac{\bar{w}}{w_L + \bar{w}}.$$

If $V_H(I) \geq (q_L + q_H)I$, this holds for all $\bar{w} \geq 0$. Otherwise it is equivalent to

$$V_H(I)(w_L + \bar{w}) \geq (q_L + q_H)I\, \bar{w} \quad \Longleftrightarrow \quad \bar{w} \leq \frac{w_L V_H(I)}{(q_L + q_H)I - V_H(I)}.$$

Combining the two inequalities and intersecting with $[w_L, w_H]$ yields (29). □

For fixed coverage I, increasing $\bar{w}$ shifts incidence from L to H exactly as under the mandate. This is the monotone incidence property used when a pooling frontier is indexed by a ceiling.

Lemma 20 (Monotone incidence for fixed coverage)**.** *Fix $I \in (0, \ell]$. On any set of feasible*

ceilings for which (28) holds, the mapping $\bar{w} \mapsto \Gamma_L(I,\bar{w}) - \Gamma_H(I,\bar{w})$ is strictly increasing on its domain.

Proof. For fixed I, $\Gamma_i(I,\bar{w}) = V_i(I) - t_i(I,\bar{w})$ and $V_i(I)$ is constant in $\bar{w}$. By (27),

$$t_H(I,\bar{w}) - t_L(I,\bar{w}) = (q_L + q_H)I\,\frac{\bar{w} - w_L}{w_L + \bar{w}}$$

is strictly increasing in $\bar{w}$, hence $t_L(I,\bar{w}) - t_H(I,\bar{w})$ is strictly decreasing and $\Gamma_L(I,\bar{w}) - \Gamma_H(I,\bar{w})$ is strictly increasing. □

Proof of Theorem 4(i). Theorem 4(i) is stated in the form needed for applications. Whenever the Pareto frontier of G^{pool} admits a one-dimensional parameterization that satisfies Assumption 2, Corollary 1 implies that the CE-indexed rule selects the Pareto-efficient pooled policy that minimizes $|\Gamma_L - \Gamma_H|$ on that frontier, equivalently minimizing worst-case veil regret.

In the ceiling-indexed pooling families used in the main text, the monotonicity requirement in Assumption 2 is supplied by Lemma 20. Lemma 19 characterizes the feasible domain of ceilings at fixed coverage, and Lemma 20 verifies strict monotonicity of $\Gamma_L - \Gamma_H$ over that domain, which is the key regularity input for the two-group reduction. □

E.5 Unraveling/separation and the fairness cost of infeasibility

When the pooling feasible set is empty, the cap-financed pooling family cannot sustain voluntary enrollment. The main text adopts the benchmark separating outcome in which group L purchases actuarially fair insurance while group H opts out.

Proof of Theorem 4(ii). Under separation and full coverage, $\Gamma_L^{\text{sep}} = V_L(\ell) - q_L\ell$ and $\Gamma_H^{\text{sep}} = 0$ by (14). Hence, using Lemma 11 (Appendix A),

$$D^{CE,\text{sep}} = |\Gamma_L^{\text{sep}} - \Gamma_H^{\text{sep}}| = V_L(\ell) - q_L\ell, \qquad S^{CE,\text{sep}} = -D^{CE,\text{sep}}.$$

If v_L is strictly concave, then $V_L(\ell) > q_L\ell$, since actuarially fair full insurance strictly improves on the risky baseline for a risk-averse agent. Therefore $D^{CE,\text{sep}} > 0$ and $S^{CE,\text{sep}} < 0$. □

E.6 Mandate versus voluntary/separation: sufficient conditions

This subsection records a convenient comparison criterion for the case in which pooling is infeasible. Let $D_*^{CE,\text{man}}$ denote the minimal attainable CE gap on the mandated menu,

$$D_*^{CE,\text{man}} := \min_{\bar{w}\in[w_L,w_H]} |\Gamma_L(\bar{w}) - \Gamma_H(\bar{w})|,$$

and let $D^{CE,\text{sep}} := V_L(\ell) - q_L\ell$ denote the separating gap.

Proposition 14 (A sufficient condition for the mandate to dominate separation)**.** *Suppose pooling is infeasible. If* $D_*^{CE,\text{man}} \leq D^{CE,\text{sep}}$*, then the mandated fair policy weakly dominates separation according to both the CE-indexed fairness criterion (smaller maximal CE gap) and the veil regret criterion (smaller worst-case regret on* $PE(G)$*).*

Proof. By Lemma 16 (Appendix D), worst-case veil regret equals D^{CE}. The CE-indexed criterion and the veil criterion therefore rank policies by the value of D^{CE} on the relevant Pareto sets. If the mandated menu achieves a gap no larger than separation, the claim follows. □

E.7 Additional comparative statics for pooling feasibility

This subsection records a sufficient condition for pooling feasibility at a fixed coverage level and its comparative statics under CARA.

Proposition 15 (A sufficient condition for pooling feasibility)**.** *Fix* $I \in (0, \ell]$*. If*

$$\underline{\bar{w}}(I) \leq \overline{\bar{w}}(I) \tag{30}$$

with $\underline{\bar{w}}(I), \overline{\bar{w}}(I)$ *as in Lemma 19, then there exists* $\bar{w} \in [w_L, w_H]$ *such that pooling with coverage* I *is feasible. Under CARA,* $\underline{\bar{w}}(I)$ *decreases in* q_L *and* a_L*, while* $\overline{\bar{w}}(I)$ *increases in* q_H *and* a_H *(when finite).*

Proof. Lemma 19 shows that feasibility for fixed I is equivalent to the existence of $\bar{w}$ in the intersection interval (29). Condition (30) ensures that $[\underline{\bar{w}}(I), \overline{\bar{w}}(I)]$ is nonempty. Under primitives for which this interval intersects $[w_L, w_H]$, a feasible ceiling exists.

Under CARA, $V_i(I)$ increases in q_i and a_i for each fixed $I > 0$ by the same log-moment argument used in Lemma 17 applied to (7). Since $\underline{\bar{w}}(I)$ is decreasing in $V_L(I)$ and $\overline{\bar{w}}(I)$ is increasing in $V_H(I)$ when finite, the stated comparative statics follow. □

Remark 4 (Scope)**.** *Proposition 15 is stated for the cap-financed proportional contribution family (27) and for CARA to keep expressions transparent. The same monotone-comparative-static logic applies to other one-dimensional incidence instruments whenever pooling feasibility is characterized by participation constraints that bound the incidence parameter and* $V_i(I)$ *shifts monotonically with risk and risk aversion.*

F Extension: incentive-compatible screening and fair insurance choice

This appendix extends the insurance application by letting feasibility be determined endogenously by incentive compatibility, participation, and budget constraints in a two-type screening environment in the spirit of Rothschild and Stiglitz (1976) (adverse selection and contract design) and Rochet (1987) (single-crossing and monotone implementability). The

purpose is to show that the CE-indexed criterion remains tractable when the feasible menu is the set of IC/IR/budget-feasible direct mechanisms rather than a family indexed by an exogenous incidence instrument. Throughout, we maintain Assumption 1 and evaluate allocations in baseline-relative CE net gains (Definition 3). Generic certainty-equivalent properties are collected in Appendix A.

F.1 Screening environment and CE net gains

There are two privately known types $i \in \{L, H\}$ with population shares $\lambda_i > 0$ and $\lambda_L + \lambda_H = 1$. Type i faces a binary loss of size $\ell > 0$ with probability $q_i \in (0, 1)$, and we maintain $q_L > q_H$. A contract specifies coverage $I \in [0, \ell]$ and a deterministic premium $p \in \mathbb{R}$. The baseline is no insurance.

Recall the CE value of coverage from Definition 8. In this environment it takes the form

$$V_i(I) := \mathrm{CE}_i(w_i,\ w_i - \ell + I) - \mathrm{CE}_i(w_i,\ w_i - \ell).$$

When deterministic premia are tracked separately, they enter additively in CE dollars under cash additivity, as in Lemma 12 (Appendix B) and Remark 1 (Appendix A). We therefore write the CE net gain from contract (I, p) as

$$\Gamma_i(I, p) := V_i(I) - p.$$

A screening policy (direct mechanism) assigns a contract to each type,

$$g = \big((I_L, p_L), (I_H, p_H)\big).$$

F.2 IC/IR and budget constraints in CE net-gain form

Write $\Gamma_i := \Gamma_i(I_i, p_i) = V_i(I_i) - p_i$ for the CE net gain delivered to type i. Define the IC wedge

$$\Delta(I) := V_L(I) - V_H(I).$$

The wedge summarizes how much more the high-risk type values coverage in CE dollars.

Assumption 7 (Single-crossing in coverage)**.** *Whenever derivatives exist,*

$$V_L'(I) > V_H'(I) \qquad \text{for all } I \in (0, \ell), \tag{31}$$

equivalently, $\Delta(\cdot)$ *is strictly increasing on* $[0, \ell]$*.*

Assumption 7 is a standard single-crossing condition for insurance screening, and it ensures that IC constraints can be expressed as bounds on the net-gain gap (Rochet, 1987, single-crossing and implementability).

Participation requires that each type weakly prefers its assigned contract to the baseline:

$$\Gamma_L \geq 0, \qquad \Gamma_H \geq 0. \tag{32}$$

Incentive compatibility is

$$\Gamma_L \geq V_L(I_H) - p_H = \Gamma_H + \Delta(I_H), \qquad \Gamma_H \geq V_H(I_L) - p_L = \Gamma_L - \Delta(I_L). \tag{33}$$

Finally, impose an ex ante no-subsidy condition (no external funds):

$$\lambda_L p_L + \lambda_H p_H \geq \lambda_L q_L I_L + \lambda_H q_H I_H. \tag{34}$$

Using $p_i = V_i(I_i) - \Gamma_i$, this becomes

$$\lambda_L \Gamma_L + \lambda_H \Gamma_H \leq \lambda_L S_L(I_L) + \lambda_H S_H(I_H), \tag{35}$$

where

$$S_i(I) := V_i(I) - q_i I \tag{36}$$

is the CE surplus net of expected cost from coverage I for type i.

Lemma 21 (Surplus properties)**.** *Maintain Assumption 1 and suppose v_i is strictly concave and continuously differentiable. Then S_i is continuous on $[0, \ell]$ and strictly increasing on $[0, \ell)$, with $S_i'(\ell) = 0$ whenever the left derivative exists. In particular, S_i is maximized at full coverage $I = \ell$.*

Proof. Continuity follows from continuity of V_i (Proposition 5 in Appendix A). Fix $I \in (0, \ell)$. Proposition 6 in Appendix A implies that V_i is differentiable on $(0, \ell)$ with

$$V_i'(I) = q_i \, \frac{v_i'(w_i - \ell + I)}{v_i'(\mathrm{CE}_i(I))}, \qquad \mathrm{CE}_i(I) := \mathrm{CE}_i(w_i,\ w_i - \ell + I).$$

The insured lottery at coverage $I < \ell$ equals $w_i - \ell + I$ in the loss state and strictly exceeds it in the no-loss state, so strict monotonicity of v_i implies

$$(1 - q_i) v_i(w_i) + q_i v_i(w_i - \ell + I) > v_i(w_i - \ell + I),$$

hence $\mathrm{CE}_i(I) > w_i - \ell + I$. Strict concavity implies v_i' is strictly decreasing, so $v_i'(\mathrm{CE}_i(I)) < v_i'(w_i - \ell + I)$ and therefore $V_i'(I) > q_i$. It follows that

$$S_i'(I) = V_i'(I) - q_i > 0 \qquad \text{for all } I \in (0, \ell).$$

At full coverage the lottery is constant at w_i, so $\mathrm{CE}_i(\ell) = w_i$ and the derivative formula yields $V_i'(\ell) = q_i$ whenever the left derivative exists, hence $S_i'(\ell) = 0$. □

F.3 Implementable set and an IC reduction

It is convenient to describe screening allocations directly in terms of

$$(I_L, I_H, \Gamma_L, \Gamma_H) \in [0, \ell]^2 \times \mathbb{R}_+^2$$

satisfying IR (32), IC (33), and budget (35).

Assumption 8 (Monotone coverage)**.** *Screening policies satisfy* $I_L \geq I_H$.

Assumption 8 is standard under single crossing and restricts attention to the Pareto-relevant region in which the higher-risk type receives weakly more coverage.

Lemma 22 (IC as bounds on the CE net-gain gap)**.** *Maintain Assumptions 7 and 8. For any* $(I_L, I_H, \Gamma_L, \Gamma_H)$*, the IC constraints* (33) *are equivalent to*

$$\Delta(I_H) \leq \Gamma_L - \Gamma_H \leq \Delta(I_L).$$

Proof. Rearranging (33) yields $\Gamma_L - \Gamma_H \geq \Delta(I_H)$ and $\Gamma_L - \Gamma_H \leq \Delta(I_L)$. Under Assumption 8 and strict increase of Δ (Assumption 7), the interval $[\Delta(I_H), \Delta(I_L)]$ is nonempty, so the two inequalities are jointly feasible and equivalent to IC. □

Let $\mathcal{F}^{IC}$ denote the set of tuples $(I_L, I_H, \Gamma_L, \Gamma_H)$ satisfying IR (32), IC (33), budget (35), and monotonicity $I_L \geq I_H$.

F.4 Pareto frontier under IC/IR/budget constraints

We study Pareto efficiency with respect to the componentwise order on CE net gains (Γ_L, Γ_H) within $\mathcal{F}^{IC}$.

Proposition 16 (Structure of Pareto-efficient screening allocations)**.** *Maintain Assumption 1 with* v_i *continuously differentiable and strictly concave, and maintain Assumptions 7 and 8. Consider Pareto efficiency in* (Γ_L, Γ_H) *within* $\mathcal{F}^{IC}$*. Then every Pareto-efficient allocation satisfies:*

1. *Budget binds:* (35) *holds with equality.*
2. *High-risk coverage is undistorted:* $I_L = \ell$.
3. *If* $I_H < \ell$ *and* $\Gamma_H > 0$*, then* IC_L *binds:* $\Gamma_L - \Gamma_H = \Delta(I_H)$.

Proof. Suppose (35) is slack at some feasible allocation. For sufficiently small $\varepsilon > 0$, the allocation $(I_L, I_H, \Gamma_L + \varepsilon, \Gamma_H + \varepsilon)$ remains feasible. IR is preserved, IC is preserved because it depends only on $\Gamma_L - \Gamma_H$ (Lemma 22), and budget feasibility holds by slackness. This is a strict Pareto improvement, so budget must bind at any Pareto-efficient allocation.

Next suppose $I_L < \ell$ at a feasible allocation. Lemma 21 implies that increasing I_L by a small $\delta > 0$ strictly increases $S_L(I_L)$ and therefore strictly increases the right-hand

side of (35). IC remains feasible because increasing I_L weakly increases the upper bound $\Delta(I_L)$ in Lemma 22, and all other constraints are unchanged. Budget slack then appears, contradicting Pareto efficiency by the previous paragraph. Therefore $I_L = \ell$ at Pareto efficiency.

Finally fix a Pareto-efficient allocation with $I_L = \ell$, $I_H < \ell$, and $\Gamma_H > 0$, and suppose IC_L is slack, so that $\Gamma_L - \Gamma_H > \Delta(I_H)$. By continuity and strict increase of Δ, there exists $\delta > 0$ with $I'_H := I_H + \delta \leq \ell$ and $\Gamma_L - \Gamma_H > \Delta(I'_H)$. The allocation $(\ell, I'_H, \Gamma_L, \Gamma_H)$ remains feasible for IR and IC, and Lemma 21 gives $S_H(I'_H) > S_H(I_H)$, so the right-hand side of (35) strictly increases and budget becomes slack. The first paragraph then yields a strict Pareto improvement, contradicting Pareto efficiency. Hence IC_L binds whenever $I_H < \ell$ and $\Gamma_H > 0$. □

Proposition 16 yields a convenient parameterization of the interior IC frontier, defined as the region with $I_H < \ell$ and $\Gamma_H > 0$.

Corollary 3 (A convenient frontier parameterization on the interior region)**.** *Maintain the assumptions of Proposition 16. On any Pareto-efficient allocation with $I_H < \ell$ and $\Gamma_H > 0$ we have $I_L = \ell$ and binding IC_L, hence for each $I_H \in [0, \ell)$ the induced CE net gains satisfy*

$$\Gamma_L = \Gamma_H + \Delta(I_H), \qquad \Gamma_H = \lambda_L S_L(\ell) + \lambda_H S_H(I_H) - \lambda_L \Delta(I_H), \tag{37}$$

and $\Gamma_H > 0$ holds if and only if the right-hand side is positive.

Proof. Binding IC_L gives $\Gamma_L = \Gamma_H + \Delta(I_H)$. Budget equality in (35) with $I_L = \ell$ then implies

$$\lambda_L(\Gamma_H + \Delta(I_H)) + \lambda_H \Gamma_H = \lambda_L S_L(\ell) + \lambda_H S_H(I_H),$$

which rearranges to (37). The positivity condition is immediate. □

F.5 Fair screening choice on the IC frontier

The CE-indexed fair-choice rule selects, from any two-identity menu, a Pareto-efficient allocation minimizing the absolute CE net-gain gap (Corollary 1). Applied to $\mathcal{F}^{IC}$, this yields the following characterization.

Theorem 6 (Fair screening characterization)**.** *Maintain the assumptions of Proposition 16 and let PE^{IC} denote the Pareto frontier of CE net gains on $\mathcal{F}^{IC}$. Then the CE-indexed fair screening allocation solves*

$$\min_{(I_L, I_H, \Gamma_L, \Gamma_H) \in PE^{IC}} |\Gamma_L - \Gamma_H|. \tag{38}$$

Moreover, if the minimizer lies on the interior region characterized in Corollary 3, then $|\Gamma_L - \Gamma_H| = \Delta(I_H)$ at the optimum, and the fair choice among interior frontier points reduces to minimizing $\Delta(I_H)$ subject to feasibility encoded by (37) and $\Gamma_H > 0$.

Proof. The first statement is Corollary 1 applied to the menu $\mathcal{F}^{IC}$ evaluated in CE net gains. On the interior region, Proposition 16 implies binding IC_L, hence $\Gamma_L - \Gamma_H = \Delta(I_H) > 0$ and

the objective equals $\Delta(I_H)$. The feasibility restriction for interior points is exactly $\Gamma_H > 0$ in (37). □

Remark 5 (Equalization and the IC wedge)**.** *Under strict single crossing, $\Delta(I) > 0$ for every $I > 0$ since $V_i(0) = 0$ and Δ is strictly increasing. On the interior IC frontier, where IC_L binds, exact equalization $\Gamma_L = \Gamma_H$ therefore requires $I_H = 0$ unless $\Delta \equiv 0$. When equalization is infeasible, the fair screening choice implements the smallest IC wedge attainable on the Pareto frontier, measured in CE dollars.*

F.6 Comparative statics: adverse-selection severity and distortions

Directional comparative statics are driven by how primitives affect the IC wedge $\Delta(I) = V_L(I) - V_H(I)$.

Proposition 17 (Comparative statics of the IC wedge)**.** *Maintain CARA utility (4). For each fixed $I \in (0, \ell]$ and each type i:*

1. *$V_i(I)$ is strictly increasing in q_i.*
2. *$V_i(I)$ is strictly increasing in a_i.*

Consequently, for each fixed $I \in (0, \ell]$, the IC wedge $\Delta(I) = V_L(I) - V_H(I)$ is strictly increasing in q_L and a_L, strictly decreasing in q_H and a_H, and, holding other primitives fixed, strictly increasing in the risk gap $q_L - q_H$.

Proof. Under CARA, (7) gives

$$V_i(I) = \frac{1}{a_i} \ln\Big((1 - q_i) + q_i e^{a_i \ell}\Big) - \frac{1}{a_i} \ln\Big((1 - q_i) + q_i e^{a_i(\ell - I)}\Big).$$

Define $B(q, a; s) := (1 - q) + q e^{as}$ and note that

$$\frac{\partial}{\partial s} \ln B(q, a; s) = a \, \frac{q e^{as}}{(1 - q) + q e^{as}}.$$

Integrating from $s = \ell - I$ to $s = \ell$ yields the representation

$$V_i(I) = \int_{\ell - I}^{\ell} \frac{q_i e^{a_i s}}{(1 - q_i) + q_i e^{a_i s}} \, ds.$$

The integrand is strictly increasing in q_i for every $s \geq 0$ and strictly increasing in a_i for every $s > 0$, so $V_i(I)$ is strictly increasing in q_i and a_i for any $I > 0$. The comparative statics for $\Delta(I)$ follow by taking differences. □

Proposition 17 implies that a larger risk gap increases the IC wedge pointwise, which tightens the IC lower bound on compressing CE net gains along the interior frontier.

Corollary 4 (Fairness cost of adverse selection under screening)**.** *Under the assumptions of Theorem 6, the minimal attainable objective value in (38) is weakly increasing under*

pointwise increases of the IC wedge $\Delta(\cdot)$. In particular, under CARA with other primitives fixed, it is increasing in the risk gap $q_L - q_H$ and in a_L and decreasing in a_H.

Proof. On the interior frontier the objective equals $\Delta(I_H)$ by Theorem 6. A pointwise upward shift of $\Delta(\cdot)$ cannot reduce the minimum value of $\Delta(I_H)$ over any fixed feasible set of I_H values. Under CARA, the wedge increases with $q_L - q_H$ and with a_L and decreases with a_H by Proposition 17. □

Remark 6 (Interface with the main paper)**.** *In this screening extension, IC/IR/budget constraints define a feasible menu, Proposition 16 and Corollary 3 identify the Pareto-relevant frontier in CE net gains, and the CE-indexed fair choice is pinned down by the same absolute-gap minimization as in the flagship application, with the IC wedge $\Delta(\cdot)$ summarizing the fairness cost of adverse selection.*

G Wealth taxation

Proposition 18 (Evaluating a statutory wealth-tax proposal on the CE-incidence frontier)**.** *Consider two identities L (non-liable) and H (liable) and fix a statutory bracket structure for a progressive wealth tax, such as Saez and Zucman (2019b). Let G^{SZ} be a restricted menu of reforms that preserve this bracket structure and vary only an implementation parameter $\theta \in [\underline{\theta}, \overline{\theta}]$ (e.g., a uniform rate scaling, an enforcement/coverage parameter, or an administratively admissible threshold adjustment), together with a deterministic rebate rule returning revenue.*[5] *Assume:*

(a) *(*Monotone incidence*) For all $\theta' < \theta''$, $c_L(\theta'') \geq c_L(\theta')$ componentwise and $c_H(\theta'') \leq c_H(\theta')$ componentwise, with at least one strict inequality in each comparison.*

(b) *(*Efficient locus*) $\theta \mapsto g(\theta)$ parameterizes the Pareto frontier within the statutory menu, $PE(G^{SZ}) = \{g(\theta) : \theta \in [\underline{\theta}, \overline{\theta}]\}$.*

Then:

(i) The CE net gains $\Gamma_L(g(\theta))$ and $-\Gamma_H(g(\theta))$ are strictly increasing and continuous in θ, and the CE net-gain gap $\theta \mapsto \Gamma_L(g(\theta)) - \Gamma_H(g(\theta))$ is strictly increasing and continuous. Hence there is at most one $\theta^\star$ satisfying $\Gamma_L(g(\theta^\star)) = \Gamma_H(g(\theta^\star))$.

(ii) Let $g^{CE} \in F^{CE}(G^{SZ})$ be the CE-indexed choice on $PE(G^{SZ})$. Then $g^{CE} = g(\theta^\star)$ if such a $\theta^\star$ exists; otherwise g^{CE} is the boundary point $g(\underline{\theta})$ or $g(\overline{\theta})$ that minimizes the maximal CE gap $D^{CE}(g(\theta))$ (Definition 5).

(iii) Let $g^{SZ} := g(1)$ denote the policy that implements the statutory schedule at full strength within this family (normalized so that $\theta = 1$ corresponds to the proposal). Its CE

[5] Formally, each $g(\theta) \in G^{SZ}$ induces state-contingent resources $c_L(\theta), c_H(\theta) \in \mathbb{R}^M$; deterministic payments can be absorbed into $c_i(\theta)$ without loss, cf. Definition 3.

fairness shortfall *relative to the CE-indexed choice is*

$$\Delta_{SZ}^{CE} := D^{CE}(g^{SZ}) - \min_{\theta\in[\underline{\theta},\bar{\theta}]} D^{CE}(g(\theta)).$$

If $\theta^\star$ exists, then $\min_\theta D^{CE}(g(\theta)) = 0$ and therefore $\Delta_{SZ}^{CE} = D^{CE}(g^{SZ}) = |\Gamma_L(g^{SZ}) - \Gamma_H(g^{SZ})|$. Moreover, by Theorem 2, this same quantity is the increase in worst-case ex post regret (in CE dollars) from adopting g^{SZ} rather than g^{CE}.

Proof. Under Assumption 1, each certainty equivalent $\mathrm{CE}_i(\cdot)$ is strictly increasing and continuous with respect to the componentwise order on $\mathbb{R}^M$ (Lemma 1). Under (a), $\theta'' > \theta'$ implies $c_L(\theta'') \geq c_L(\theta')$ with a strict inequality in at least one state, hence $\mathrm{CE}_L(c_L(\theta'')) > \mathrm{CE}_L(c_L(\theta'))$ and therefore $\Gamma_L(g(\theta'')) > \Gamma_L(g(\theta'))$. Similarly, $c_H(\theta'') \leq c_H(\theta')$ with a strict inequality implies $\mathrm{CE}_H(c_H(\theta'')) < \mathrm{CE}_H(c_H(\theta'))$ and thus $\Gamma_H(g(\theta'')) < \Gamma_H(g(\theta'))$. Continuity follows from continuity of CE_i. The gap $\Gamma_L - \Gamma_H$ is therefore strictly increasing and continuous, proving (i). Statement (ii) is exactly the two-group characterization of the CE-indexed rule applied to $\mathrm{PE}(G^{SZ})$ (Proposition 4). For (iii), the definition of Δ_{SZ}^{CE} is immediate, and when an equalizer exists the minimum of D^{CE} over θ is 0 by Definition 5. The veil interpretation follows from Theorem 2. □

Proposition 19 turns evaluation of θ^{SZ} into a checkable first-order certificate: either θ^{SZ} satisfies (40) (CE-indexed optimal within the statutory class) or there is a statutory direction that strictly reduces the maximal CE complaint.

Proposition 19 (A general optimality certificate for statutory wealth-tax proposals)**.** *Maintain Assumption 1. Fix a statutory wealth-tax bracket structure (e.g., Saez and Zucman, 2019b) and let $\Theta \subset \mathbb{R}^K$ be a nonempty compact convex set of* implementable *parameters that preserve this structure (rates, thresholds, enforcement/coverage parameters, valuation rules, etc.), together with a deterministic rebate rule that balances the budget. Let $g : \Theta \to PE(G^{SZ})$ parameterize the Pareto frontier* within the statutory menu*, and write*

$$\Gamma_i(\theta) := \Gamma_i(g(\theta)), \qquad D^{CE}(\theta) := \max_{i\in N} \Gamma_i(\theta) - \min_{i\in N} \Gamma_i(\theta).$$

Let $\theta^{SZ} \in \Theta$ denote the parameter that implements the Saez–Zucman proposal within this class.[6]

Assume that each $\Gamma_i(\cdot)$ is differentiable on a neighborhood of Θ.[7] *Define the active sets*

$$I^+(\theta) := \arg\max_{i\in N} \Gamma_i(\theta), \qquad I^-(\theta) := \arg\min_{i\in N} \Gamma_i(\theta),$$

and let $N_\Theta(\theta)$ denote the normal cone to Θ at θ.

*(i) (*Directional derivative / envelope formula*) For any direction $d \in \mathbb{R}^K$, the directional*

[6] For instance, θ^{SZ} may encode the Warren version with fixed bracket rates and thresholds, together with the proposal's intended enforcement intensity.

[7] This is a mild regularity condition on the reduced-form mapping from implementable statutory parameters to state-contingent after-tax resources; it allows for behavioral responses and general equilibrium effects, provided the induced consumption vectors vary smoothly with θ.

derivative of D^{CE} at θ exists and equals

$$D^{CE\,\prime}(\theta; d) = \max_{i \in I^+(\theta)} \nabla\Gamma_i(\theta) \cdot d \; - \; \min_{j \in I^-(\theta)} \nabla\Gamma_j(\theta) \cdot d. \tag{39}$$

*(ii) (*First-order certificate for CE-indexed optimality*) A parameter $\theta^\star \in \Theta$ minimizes $D^{CE}(\theta)$ over Θ if and only if there exist probability vectors $u \in \Delta(I^+(\theta^\star))$ and $v \in \Delta(I^-(\theta^\star))$ and a normal vector $n \in N_\Theta(\theta^\star)$ such that*

$$\sum_{i \in I^+(\theta^\star)} u_i \, \nabla\Gamma_i(\theta^\star) \; - \; \sum_{j \in I^-(\theta^\star)} v_j \, \nabla\Gamma_j(\theta^\star) \; + \; n \; = \; 0. \tag{40}$$

Equivalently, $0 \in co\{\nabla\Gamma_i(\theta^\star) : i \in I^+(\theta^\star)\} - co\{\nabla\Gamma_j(\theta^\star) : j \in I^-(\theta^\star)\} + N_\Theta(\theta^\star)$.

*(iii) (*Rigorous evaluation of a proposal*) If (40) fails at θ^{SZ}, then there exists a feasible direction $d \in T_\Theta(\theta^{SZ})$ (tangent cone) such that $D^{CE\,\prime}(\theta^{SZ}; d) < 0$. In particular, there is a local statutory modification of the proposal—changing only implementable parameters within Θ and remaining on $PE(G^{SZ})$—that strictly reduces the maximal CE complaint.*

*(iv) (*Convex benchmark and computable near-optimality*) If, in addition, $\theta \mapsto \Gamma_i(\theta)$ is affine for each $i \in N$, then D^{CE} is convex on Θ and (40) is a global optimality certificate. Moreover, for any $\theta \in \Theta$ and any choice of $u \in \Delta(I^+(\theta))$, $v \in \Delta(I^-(\theta))$,*

$$D^{CE}(\theta) - \min_{\vartheta \in \Theta} D^{CE}(\vartheta) \; \le \; \sup_{\vartheta \in \Theta} \Big(\Big[\sum_{i \in I^+(\theta)} u_i \, \nabla\Gamma_i(\theta) - \sum_{j \in I^-(\theta)} v_j \, \nabla\Gamma_j(\theta) \Big] \cdot (\theta - \vartheta) \Big),$$

yielding a computable upper bound on the fairness shortfall of θ using only local derivatives and the geometry of Θ.

Proof. Part (i) is an envelope calculation for the max and min operators. Since each Γ_i is differentiable, the directional derivative of $U(\theta) := \max_i \Gamma_i(\theta)$ is $U'(\theta; d) = \max_{i \in I^+(\theta)} \nabla\Gamma_i(\theta) \cdot d$, and the directional derivative of $L(\theta) := \min_i \Gamma_i(\theta)$ is $L'(\theta; d) = \min_{j \in I^-(\theta)} \nabla\Gamma_j(\theta) \cdot d$. Because $D^{CE} = U - L$, (39) follows.

For (ii)–(iii), note that minimizing D^{CE} over Θ is equivalent to minimizing the convex, locally Lipschitz function $U(\theta) + (-L(\theta))$ over a compact convex set. The Clarke subdifferential of U at θ is $\partial U(\theta) = \mathrm{co}\{\nabla\Gamma_i(\theta) : i \in I^+(\theta)\}$, while the subdifferential of $-L$ is $\partial(-L)(\theta) = \mathrm{co}\{-\nabla\Gamma_j(\theta) : j \in I^-(\theta)\}$. Therefore

$$\partial D^{CE}(\theta) = \partial U(\theta) + \partial(-L)(\theta) = \mathrm{co}\{\nabla\Gamma_i(\theta) : i \in I^+(\theta)\} - \mathrm{co}\{\nabla\Gamma_j(\theta) : j \in I^-(\theta)\}.$$

A point $\theta^\star \in \Theta$ minimizes D^{CE} over Θ if and only if $0 \in \partial D^{CE}(\theta^\star) + N_\Theta(\theta^\star)$, which is exactly (40) (writing a subgradient as convex combinations over the active sets), proving (ii). If (40) fails at θ^{SZ}, then $0 \notin \partial D^{CE}(\theta^{SZ}) + N_\Theta(\theta^{SZ})$, and the separation theorem implies the existence of a feasible direction $d \in T_\Theta(\theta^{SZ})$ with strictly negative directional derivative, yielding (iii) via (i).

For (iv), if each Γ_i is affine, then $U(\theta) = \max_i \Gamma_i(\theta)$ and $-L(\theta) = \max_i(-\Gamma_i(\theta))$ are convex, hence D^{CE} is convex. In that case, the subgradient optimality condition in (ii) is

sufficient for global optimality. The bound is the standard convexity inequality $D^{CE}(\vartheta) \geq D^{CE}(\theta) + s \cdot (\vartheta - \theta)$ for any subgradient $s \in \partial D^{CE}(\theta)$, rearranged and maximized over $\vartheta \in \Theta$. □

H Existence, continuity, and stability of CE-indexed selections

This appendix records existence and stability results that are used whenever the CE-indexed fair-choice rule is applied to a parameterized family of constrained menus. For each parameter value, the Pareto set is nonempty and compact, the CE-indexed continuation program admits minimizers, and the resulting correspondences vary stably with primitives under standard continuity conditions.

We use attainment on compact sets and continuity of max-gap functionals from Appendix A. We also use that Pareto comparisons in consumption space coincide with componentwise comparisons of CE net gains (Lemma 13 in Appendix C).

H.1 A general parametric menu environment

Let (Θ, d_Θ) be a compact metric parameter space. Let $(\mathcal{G}, d_\mathcal{G})$ be a compact metric space of policies, and let $G : \Theta \rightrightarrows \mathcal{G}$ be a nonempty-valued correspondence, where $G(\theta)$ is the feasible policy menu at parameter θ.

Each pair (g, θ) induces for each identity $i \in N$ a state-contingent consumption vector $c_i(g, \theta) \in \mathbb{R}^M$. To keep notation light, deterministic payments are absorbed into $c_i(g, \theta)$ in this appendix. When deterministic payments are separated explicitly, the only additional requirement for translating them outside the CE operator is the cash-additivity condition in Lemma 12 (Appendix B) and Remark 1 (Appendix A).

Fix a baseline selection $g^0(\theta) \in G(\theta)$ and define baseline-relative CE net gains by

$$\Gamma_i(g, \theta) := \mathrm{CE}_i\big(c_i(g, \theta)\big) - \mathrm{CE}_i\big(c_i(g^0(\theta), \theta)\big),$$

where CE_i is defined under Assumption 1.

Assumption 9 (Primitive regularity)**.**

1. Menu compactness and Hausdorff continuity. *For each θ, $G(\theta)$ is nonempty and compact. If $\theta_n \to \theta$, then $d_H\big(G(\theta_n), G(\theta)\big) \to 0$.*
2. Outcome continuity. *For each $i \in N$, $(g, \theta) \mapsto c_i(g, \theta)$ is continuous on $\mathcal{G} \times \Theta$.*
3. Baseline continuity. *The baseline selection $\theta \mapsto g^0(\theta)$ is continuous.*
4. CE well-definedness. *Assumption 1 holds for each identity, so CE_i is well-defined and continuous on its admissible domain (Lemma 6 in Appendix A).*
5. No disappearance of efficient policies. *If $\theta_n \to \theta$ and $g \in PE(G(\theta))$, then there exist $g_n \in PE(G(\theta_n))$ with $g_n \to g$.*

Assumption 9(5) is a lower hemicontinuity condition for the Pareto correspondence. It holds automatically in the one-dimensional frontier environments that drive the applications in the main text, and it is the condition needed for stability of continuation selections when the efficient set is not single-valued.

Lemma 23 (Continuity of CE net gains)**.** *Maintain Assumption 9. For each $i \in N$, the mapping $(g,\theta) \mapsto \Gamma_i(g,\theta)$ is continuous on $\mathcal{G} \times \Theta$. Consequently, for each θ the maximal CE gap*

$$D^{CE}(g,\theta) := \max_{i,j \in N} \big(\Gamma_j(g,\theta) - \Gamma_i(g,\theta)\big)$$

is continuous on $G(\theta)$ (and hence on $PE(G(\theta))$).

Proof. By Assumption 9(2) and (3), both $(g,\theta) \mapsto c_i(g,\theta)$ and $(g,\theta) \mapsto c_i(g^0(\theta),\theta)$ are continuous. Continuity of CE $_i$ (Lemma 6 in Appendix A) and composition yield continuity of $(g,\theta) \mapsto \Gamma_i(g,\theta)$. The map $D^{CE}(\cdot,\theta)$ is a finite maximum of continuous functions and is therefore continuous by Lemma 10 (Appendix A). □

H.2 Pareto sets and continuation sets

For each θ, define the CE-image of the menu by

$$\Gamma(G(\theta)) := \{\Gamma(g,\theta) \in \mathbb{R}^N : \ g \in G(\theta)\}.$$

By Assumption 9 and Lemma 23, $\Gamma(G(\theta))$ is compact for each θ.

Definition 12 (CE-maximal set)**.** *For a compact set $K \subseteq \mathbb{R}^N$, define the set of maximal elements under the componentwise order by*

$$Max(K) := \Big\{x \in K : \ \nexists\, y \in K \text{ with } y \geq x \text{ and } y \neq x\Big\}.$$

Lemma 24 (Nonemptiness and compactness of Pareto sets)**.** *Maintain Assumption 9. For each $\theta \in \Theta$:*

1. *$Max(\Gamma(G(\theta)))$ is nonempty and compact.*
2. *The Pareto set satisfies*

$$PE(G(\theta)) = \{g \in G(\theta) : \ \Gamma(g,\theta) \in Max(\Gamma(G(\theta)))\}.$$

In particular, $PE(G(\theta))$ is nonempty and compact.

Proof. (1) Let $K := \Gamma(G(\theta))$, which is nonempty and compact. The continuous scalarization $x \mapsto \sum_{i \in N} x_i$ attains a maximum on K (Lemma 9 in Appendix A). Any maximizer is maximal under the product order, since otherwise one could strictly increase the sum, so Max (K) is nonempty. Compactness follows because Max (K) is closed in K.

(2) By Lemma 13 (Appendix C), Pareto dominance in $G(\theta)$ is equivalent to componentwise dominance in $\Gamma(G(\theta))$. Hence g is Pareto efficient in $G(\theta)$ if and only if $\Gamma(g,\theta)$ is maximal in $\Gamma(G(\theta))$. □

H.3 Existence of CE-indexed continuation selections

Define the CE continuation objective on $\mathrm{PE}\,(G(\theta))$ by

$$D^{CE}(g,\theta) = \max_{i,j\in N}\big(\Gamma_j(g,\theta) - \Gamma_i(g,\theta)\big), \qquad S^{CE}(g,\theta) = -D^{CE}(g,\theta),$$

and define the CE-indexed continuation correspondence by

$$\mathcal{S}^{CE}(\theta) := \arg\min_{g\in\mathrm{PE}\,(G(\theta))} D^{CE}(g,\theta) = \arg\max_{g\in\mathrm{PE}\,(G(\theta))} S^{CE}(g,\theta).$$

Proposition 20 (Existence of minimizers)**.** *Maintain Assumption 9. For every $\theta \in \Theta$, $\mathcal{S}^{CE}(\theta)$ is nonempty and compact.*

Proof. By Lemma 24, $\mathrm{PE}\,(G(\theta))$ is nonempty and compact. By Lemma 23, $D^{CE}(\cdot,\theta)$ is continuous on $\mathrm{PE}\,(G(\theta))$. Existence and compactness of argmin sets follow from Lemma 9 (Appendix A). □

H.4 Stability under perturbations

This subsection records stability of the Pareto correspondence and of the CE continuation correspondence. Upper hemicontinuity is the appropriate notion when the argmin set need not be single-valued.

Lemma 25 (Upper hemicontinuity of Pareto sets)**.** *Maintain Assumption 9. The correspondence $\theta \mapsto PE\,(G(\theta))$ is upper hemicontinuous with nonempty compact values.*

Proof. Nonempty compact values follow from Lemma 24. It is enough to show that the graph of $\mathrm{PE}\,(\cdot)$ is closed. Let $\theta_n \to \theta$ and $g_n \to g$ with $g_n \in \mathrm{PE}\,(G(\theta_n))$. Hausdorff continuity and $g_n \in G(\theta_n)$ imply $g \in G(\theta)$.

Suppose for contradiction that $g \notin \mathrm{PE}\,(G(\theta))$. Then there exists $h \in G(\theta)$ such that $\Gamma(h,\theta) \geq \Gamma(g,\theta)$ componentwise with strict inequality for some identity. Hausdorff continuity of G yields $h_n \in G(\theta_n)$ with $h_n \to h$. Continuity of Γ (Lemma 23) implies that for all large n one has $\Gamma(h_n,\theta_n) \geq \Gamma(g_n,\theta_n)$ componentwise with strict inequality for some identity. Lemma 13 (Appendix C) then implies that h_n Pareto dominates g_n in $G(\theta_n)$, contradicting $g_n \in \mathrm{PE}\,(G(\theta_n))$. Hence $g \in \mathrm{PE}\,(G(\theta))$ and the graph is closed. □

Proposition 21 (Upper hemicontinuity of CE continuation selections)**.** *Maintain Assumption 9. The correspondence $\theta \mapsto \mathcal{S}^{CE}(\theta)$ is upper hemicontinuous with nonempty compact values.*

Proof. Nonempty compact values follow from Proposition 20. It is enough to show that the graph of $\mathcal{S}^{CE}(\cdot)$ is closed. Let $\theta_n \to \theta$ and $g_n \to g$ with $g_n \in \mathcal{S}^{CE}(\theta_n)$.

Since $g_n \in \mathcal{S}^{CE}(\theta_n)$ implies $g_n \in \mathrm{PE}\,(G(\theta_n))$, Lemma 25 yields $g \in \mathrm{PE}\,(G(\theta))$. Suppose for contradiction that $g \notin \mathcal{S}^{CE}(\theta)$. Then there exists $h \in \mathrm{PE}\,(G(\theta))$ and $\varepsilon > 0$ such that

$$D^{CE}(g, \theta) \geq D^{CE}(h, \theta) + \varepsilon.$$

Assumption 9(5) provides $h_n \in \mathrm{PE}\,(G(\theta_n))$ with $h_n \to h$. Continuity of D^{CE} (Lemma 23) implies $D^{CE}(g_n, \theta_n) \to D^{CE}(g, \theta)$ and $D^{CE}(h_n, \theta_n) \to D^{CE}(h, \theta)$. For large n we therefore have

$$D^{CE}(g_n, \theta_n) \geq D^{CE}(h_n, \theta_n) + \varepsilon/2.$$

This contradicts that g_n minimizes $D^{CE}(\cdot, \theta_n)$ over $\mathrm{PE}\,(G(\theta_n))$. Hence $g \in \mathcal{S}^{CE}(\theta)$ and the graph is closed. □

Corollary 5 (Continuity under uniqueness)**.** *Maintain Assumption 9. If $\mathcal{S}^{CE}(\theta) = \{g^*(\theta)\}$ is single-valued at θ, then the selection function $\theta \mapsto g^*(\theta)$ is continuous at θ.*

Proof. Let $\theta_n \to \theta$ and choose $g_n \in \mathcal{S}^{CE}(\theta_n)$. Compactness of $\mathcal{G}$ implies (g_n) has cluster points. Upper hemicontinuity (Proposition 21) implies every cluster point lies in $\mathcal{S}^{CE}(\theta) = \{g^*(\theta)\}$, hence $g_n \to g^*(\theta)$. □

H.5 A uniqueness route in the two-identity case

The refined selector in the companion paper provides a canonical tie-break and generic uniqueness under its regularity conditions. In the two-identity applications that drive Section 7, uniqueness can often be verified directly from strict monotonicity of the CE net-gain gap along a one-dimensional efficient frontier.

Lemma 26 (Two identities: uniqueness from strict monotonicity along a one-dimensional frontier)**.** *Suppose $N = \{1, 2\}$ and $PE(G(\theta))$ is parameterized by $\tau \in [\underline{\tau}, \overline{\tau}]$ via $g(\tau)$ such that $\tau \mapsto \Gamma_1(g(\tau), \theta) - \Gamma_2(g(\tau), \theta)$ is continuous and strictly monotone. Then $\mathcal{S}^{CE}(\theta)$ is single-valued.*

Proof. By Lemma 15 (Appendix C), $D^{CE}(g(\tau), \theta) = |\Gamma_1(g(\tau), \theta) - \Gamma_2(g(\tau), \theta)|$. If the difference crosses zero, strict monotonicity implies a unique root and hence a unique minimizer with $D^{CE} = 0$. If it does not cross zero, strict monotonicity implies the absolute value is minimized at the unique boundary point closest to zero. In both cases the minimizer is unique. □

Remark 7 (Interface with the refined selector F^{CE})**.** *Appendix C identifies the CE-indexed specialization of the companion selector with the explicit program defining $\mathcal{S}^{CE}(\theta)$, up to the shared refinement. Proposition 21 and Corollary 5 therefore deliver stability of the chosen policy under perturbations of primitives, and Lemma 26 provides a simple uniqueness route in two-identity incidence-frontier environments.*