# Online Appendix to

# Rights Transitions in Dynamic Assignment: Safety, Investment, and Future Exchange

Dalin Sheng*

August 2026

This Online Appendix contains supplementary characterizations, examples, and complete proofs for results stated in the main article. Notation and theorem references not redefined here refer to the main article.

## A Supplementary Results on Support-Safe Transitions

This section collects the supplementary characterizations referenced in the main article. General rights changes admit a core and blocking-cycle test and an exact, exhaustive optimization procedure. For pure settlement, benchmark screening isolates the only difficult case and supplies the proof inputs for Proposition 3; a closed-form frontier covers first-round cycle vesting and its recursive extension; and a regret representation shows that safety is exactly worst-case no-harm.

The signature has a core-theoretic representation. For $a \in \mathcal{A}_F^H$, let $y^a = \Phi_F^g(P_F^a)$ be the benchmark allocation and let $\bar{q}$ extend $q$ by assigning each entrant her entry object. Construct the directed graph $D^a(q)$ on $N \cup E(F)$ with an arc $i \to j$ whenever $i$ weakly prefers $\bar{q}_j$ to $y_i^a$ at the probe profile. Mark the arc strict when the preference is strict.

**Theorem OA.1** (Common-core support and blocking cycles)**.** *A rights state $q$ is $F$-safe from $g$ if and only if, for every $a \in \mathcal{A}_F^H$, the benchmark allocation $y^a$ is the unique strong-core allocation under rights $\bar{q}$. Equivalently, for every $a$, the graph $D^a(q)$ has no directed cycle containing a strict arc. Equivalently again, no strict arc has both endpoints in the same strongly connected component.*

*Southwestern University of Finance and Economics. Email: shengdl@swufe.edu.cn. This work was supported by the National Natural Science Foundation of China (grant number 72503198).

*Proof.* With strict preferences, TTC selects the unique strong-core allocation of a Shapley–Scarf market (Roth and Postlewaite, 1977). Theorem 1 therefore gives the first equivalence. A directed cycle in $D^a(q)$ assigns each member the $\bar{q}$-endowment of her successor. It weakly blocks $y^a$ exactly when every arc is weak and at least one is strict. Conversely, a weakly blocking coalition's permutation decomposes into directed cycles, at least one of which contains a strict improvement. Finally, a directed cycle contains a given strict arc if and only if that arc's endpoints lie in the same strongly connected component. □

**Corollary OA.2** (History-dependent transitions)**.** *A nonanticipatory rights transition is support-safe after every history if and only if, at each history, the state it selects belongs to* $\mathcal{Q}_{F(\mathfrak{h})}(g(\mathfrak{h}))$ *for that history's current rights and announced finite support.*

*Proof.* Condition on a history and apply Theorem 1. The transition is fixed before the realized extension by nonanticipation. The characterization permits arbitrary dependence on the observed history. □

Testing supported objects one at a time is not equivalent to testing their joint arrival. The following example isolates the gap.

**Example OA.1** (Singleton probes do not certify a batch)**.** Let $N = \{1, 2, 3\}$, $H = \{a, b, c\}$, $g = (a, b, c)$, and $q = (c, a, b)$. Future support is $F = \{u, v\}$, and incumbent preferences are

$$\begin{aligned} &1\text{: } b \succ c \succ a \succ v \succ u, \\ &2\text{: } v \succ c \succ u \succ a \succ b, \\ &3\text{: } a \succ v \succ u \succ c \succ b. \end{aligned}$$

For a singleton entrant, the probe outcomes under $g$ and $q$ coincide as follows; allocations list incumbents first and the entrant last.

| entry object | probe target | common TTC allocation |
|---|---|---|
| $u$ | $a, b, c,$ or $u$ | $(b, c, a, u)$ |
| $v$ | $a$ | $(b, v, c, a)$ |
| $v$ | $b$ | $(c, v, a, b)$ |
| $v$ | $c$ | $(b, v, a, c)$ |
| $v$ | $v$ | $(b, c, a, v)$ |

Lemma OA.13 therefore implies neutrality to every preference of either singleton entrant. If both arrive, let the owner of $u$ target $b$ and the owner of $v$ target $a$. The elementary probe profile gives

$$\Phi^g_{\{u,v\}} = (b, v, c, u, a) \quad \text{and} \quad \Phi^q_{\{u,v\}} = (c, v, u, b, a).$$

Thus $q$ is safe for each singleton support, yet unsafe for their union.

The characterization is operational. It also defines a maximum safe transition over all bijective rights states. Let $x$ be the current implemented allocation and let $V_x(q) = \{i \in N : q_i = x_i\}$.

**Proposition OA.3** (Maximum safe transition)**.** *There exists a maximum-vesting safe transition*

$$q^\star \in \arg\max_{q \in \mathcal{Q}_F(g)} |V_x(q)|.$$

*It can be computed exactly by comparing probe signatures over all $n!$ bijective rights states. If $n = |N|$ and $m = |F|$, the maximal-batch reduction and entrant-path compression leave $(n+1)^m$ probe profiles per rights state. A standard quadratic TTC implementation therefore gives running time*

$$O\big(n!\,(n+1)^m(n+m)^2\big)\,.$$

*For fixed $m$, checking whether a given transition is safe is polynomial in $n$.*

*If the admissible transitions are restricted only for implementation purposes to $q_i \in \{g_i, x_i\}$, let $\pi(i) = g^{-1}(x_i)$. A selected set of agents can vest and preserve bijectivity if and only if it is a union of cycles of $\pi$. If $\pi$ has $c$ nontrivial cycles, a maximum safe transition in this pure-settlement class can therefore be found by the baseline procedure of testing all $2^c$ cycle collections, in time $O(2^c(n+1)^m(n+m)^2)$. This brute-force specialization is exact across all TTC rounds and all joint cycle selections. The main article's Theorem 2 proves that safety factorizes across those cycles and removes the $2^c$ factor.*

*Proof.* The safe class is nonempty because it contains $g$, and it is finite, so a maximizer exists. Theorem 1 proves that the signature comparison accepts exactly the safe states. In the maximal batch, each of the $m$ entrants need only target one of the $n$ incumbent objects or keep her own entry object. Lemmas OA.14 and OA.15 make those $(n+1)^m$ profiles sufficient, giving the displayed running time. The signatures may be compared sequentially, so no exponential-size signature need be stored.

For the specialization, write $q_i = g_{\rho(i)}$, where $\rho(i) = \pi(i)$ if $i$ vests and $\rho(i) = i$ otherwise. For each $j$, the number of preimages under $\rho$ is

$$\mathbf{1}\{j \text{ does not vest}\} + \mathbf{1}\{\pi^{-1}(j) \text{ vests}\}.$$

This number is one for every $j$ if and only if vesting is constant along each cycle of $\pi$. Hence the admissible sets are exactly the $2^c$ unions of nontrivial cycles. Testing their signatures and selecting one with maximum vested cardinality proves the claim. □

Maximum cardinality need not select a unique rights state; an announced lexicographic tie-break can select one representative. More general institutional priorities can be represented by maximizing $\sum_i w_{i,q_i}$ over the same exact safe class. The factorial term comes from the exact enumeration used here to optimize over arbitrary bijections; verifying a proposed transition, by contrast, remains polynomial for fixed $m$.

## A.1 Benchmark Screening for Pure Settlement

For pure settlement, part of the comparison between two rights states can be read from the original-rights benchmark alone. Relative to a selected member $k \in V(\mathcal{S})$, call the probe *neutral*, *tempting*, or *wounding* according as

$$y_k^a = x_k, \qquad y_k^a \succ_k x_k, \qquad \text{or} \qquad x_k \succ_k y_k^a.$$

The terminology separates two pointwise implications that go in opposite directions.

**Proposition OA.4** (Benchmark screening)**.** *Fix $\mathcal{S} \subseteq \mathcal{C}$ and a probe $a$.*

*(i) If the probe wounds some $k \in V(\mathcal{S})$, then it separates $g^{\mathcal{S}}$ from $g$.*

*(ii) If the probe tempts no member of $V(\mathcal{S})$, then it does not separate $g^{\mathcal{S}}$ from $g$.*

*(iii) If no elementary probe ever tempts a member of $V(\mathcal{S})$, then $g^{\mathcal{S}}$ is $F$-safe from $g$. In particular, every union of never-tempted cycles is safe.*

*Proof.* For (i), a selected agent $k$ owns $x_k$ under $g^{\mathcal{S}}$. Individual rationality gives

$$\Phi_F^{g^{\mathcal{S}}}(P_F^a)_k \succeq_k x_k \succ_k y_k^a,$$

so the allocations differ.

For (ii), put $q = g^{\mathcal{S}}$ and $y = y^a$. Extend $\pi$ by the identity off $V(\mathcal{S})$ and define

$$\tau(v) = \begin{cases} \pi(v), & v \in V(\mathcal{S}), \\ v, & v \notin V(\mathcal{S}). \end{cases}$$

On the maximal-batch market, the extended rights satisfy $\bar{q}_v = \bar{g}_{\tau(v)}$. Moreover,

$$\bar{g}_{\pi(k)} = x_k \succeq_k y_k \qquad \text{for every } k \in V(\mathcal{S}), \tag{1}$$

where the inequality is exactly the absence of temptation.

Suppose $y$ were not the TTC allocation under $\bar{q}$. By Theorem OA.1, a weakly blocking cycle under $\bar{q}$ would exist, with at least one strict arc. Replace every arc $u \to v$ in that

cycle by $u \to \tau(v)$; the object offered to $u$ is unchanged. Whenever $v \in V(\mathcal{S})$, append the settlement path

$$\pi(v) \to \pi^2(v) \to \cdots \to v.$$

Every appended arc is valid by (1). The result is a closed directed walk under $\bar{g}$ that retains a strict arc. Decomposing it into directed cycles produces a weakly blocking cycle against $y$ under $\bar{g}$, contradicting that $y$ is the original-rights TTC allocation. Thus the probe allocations coincide. Applying (ii) at every elementary probe and then Theorem 1 proves (iii). □

**Proposition OA.5** (Vested-neutral reverse screening)**.** *Fix $\mathcal{S} \subseteq \mathcal{C}$ and an elementary probe $a$, and put $q = g^{\mathcal{S}}$ and $\widehat{y}^a = \Phi_F^q(P_F^a)$. If*

$$\widehat{y}_k^a = x_k \quad \text{and} \quad y_k^a \succeq_k x_k \qquad \text{for every } k \in V(\mathcal{S}),$$

*then $\widehat{y}^a = y^a$. Equivalently, if vested rights leave every selected member at her vested object and change the allocation, the original-rights benchmark wounds a selected member at that probe.*

*Proof.* Suppress the probe superscript. Original-rights individual rationality gives $y_i \succeq_i g_i = q_i$ for every agent outside $V(\mathcal{S})$; the displayed hypothesis gives $y_k \succeq_k x_k = q_k$ inside it. Hence $y$ is individually rational under $q$.

Execute TTC under $q$ one cycle at a time. If $y \neq \widehat{y}$, let $L$ be the first executed cycle whose assignments do not all agree with $y$. Since all earlier cycles receive their $y$-objects, every $y_i$ for $i \in L$ is still available. The TTC object assigned along $L$ is therefore weakly preferred to $y_i$ by every member of $L$, with at least one strict comparison. Thus $L$ weakly blocks $y$ under $q$.

No selected member can lie in $L$. If selected $k$ belongs to a TTC cycle under $q$, her final assignment $\widehat{y}_k = q_k$ forces that cycle to be the singleton $\{k\}$. A discrepant singleton would have $q_k \succ_k y_k$, contrary to the hypothesis. Hence $L$ contains only agents whose rights are unchanged, so $q_i = g_i$ for every $i \in L$. The same cycle then weakly blocks the original-rights TTC allocation $y$ under $g$, a contradiction to Theorem OA.1. □

Both propositions are pointwise: they settle every benchmark with no selected temptation and every benchmark with a selected wound. The difficult case in Proposition 3 is a selected collection that is tempted somewhere but wounded nowhere; Section A.4 constructs a separator for that case.

## A.2 A Closed-Form Frontier for First-Round Cycle Vesting

The probe signature is a combinatorial test for arbitrary rights transitions. For the common policy of vesting complete cycles of the current TTC allocation, however, the first TTC round admits a local upper-contour representation. Continue with the main article's market $M$ and support $F$.

Let $C$ be a nontrivial first-round TTC cycle in $M$. Write $x_i$ for the object assigned to $i \in C$. Because the cycle occurs in the first round, $x_i$ is $i$'s favorite object in $H$. Vesting $C$ changes rights to $g^C$, where $g_i^C = x_i$ for $i \in C$ and all other rights are unchanged. For a collection $\mathcal{S}$ of first-round cycles, let $g^{\mathcal{S}}$ vest every cycle in $\mathcal{S}$.

For a feasible extension $E$ and incumbent rights state $q$, write $\mathrm{TTC}(E; q)$ for final TTC from $q$ together with the entrants' entry rights.

**Definition OA.1** (Extension safety and future satiation)**.** Cycle $C$ is *extension-safe relative to* $F$ if

$$\mathrm{TTC}(E; g) = \mathrm{TTC}(E; g^C) \quad \text{for every extension } E \text{ drawn from } F.$$

It is *future-satiated relative to* $F$ if

$$x_i \succ_i z \quad \text{for every } i \in C \text{ and every } z \in F.$$

A collection $\mathcal{S}$ is extension-safe if $\mathrm{TTC}(E; g) = \mathrm{TTC}(E; g^{\mathcal{S}})$ for every such extension.

Future satiation says that each member receives her favorite object in the entire feasible universe $H \cup F$. It uses only upper-contour information; the institution need not know which entrants arrive or how they rank objects.

A future-satiated cycle is absorbing. Its members make the same choices whether the cycle remains provisional or vests, and the same residual market remains after they leave. Failure of satiation creates a probe: a tempting future object can enter with an owner who demands the exposed member's pre-vesting right. These two observations give the section's support-dependent characterization.

**Theorem OA.6** (Exact frontier for first-round cycle vesting)**.** *The following statements hold relative to* $F$*.*

*(i) A nontrivial first-round TTC cycle is extension-safe if and only if it is future-satiated.*

*(ii) A collection of first-round cycles is extension-safe if and only if every selected nontrivial cycle is future-satiated.*

*Consequently, up to vesting trivial cycles (which leaves rights unchanged), the unique inclusion-maximal safe collection consists of all future-satiated nontrivial first-round cycles.*

*Proof.* The proof separates absorption from exposure. For sufficiency, suppose first that $C$ is future-satiated. In every feasible extension, each $i \in C$ continues to point to $x_i$. Under original rights, the members form the same first-round cycle and receive their current assignments. Under $g^C$, each member points to and owns $x_i$, so the members form self-cycles. Both histories remove the same agents with the same objects. All residual rights are identical. Agents outside $C$ that point into $C$ are paths into a cycle and remain after it is removed. The residual TTC problems are therefore identical, proving sufficiency. The argument applies simultaneously to disjoint safe cycles.

For exposure, suppose $z \succ_i x_i$ for some $i \in C$ and $z \in F$. Let $h = g_i$. Since $C$ is nontrivial, a distinct predecessor $p \in C$ receives $h$ in the current cycle, so $g_p^C = h$. Add only the entrant endowed with $z$ and let that entrant rank

$$h \succ_0 z \succ_0 H \setminus \{h\}.$$

Order the remaining old objects arbitrarily. Under original rights, $i$ points to $z$, entrant 0 points to $h$, and $i$ owns $h$. They form a first-round cycle, so $i$ receives $z$.

Under $g^C$, $p$ owns $h = x_p$. If $z \succ_p h$, then $p$ and the entrant form the corresponding cycle, so $p$, not $i$, receives $z$. If $h \succ_p z$, then $p$ self-cycles with $h$; after $h$ is removed, the entrant self-cycles with $z$. Again $i$ does not receive $z$. Thus the two final allocations differ. This proves necessity for one cycle. For a collection, absorption applies simultaneously to all safe selected cycles. If a selected cycle is unsafe, the same probe targets it; all other selected cycles are disjoint in agents and objects and can at most point into the displayed cycle. □

The theorem converts knowledge of the object support into a local upper-contour test. Its only input is the ordinal comparison of $F$ with current assignments, information available before any entrant arrives. Within first-round cycle vesting, early finality is safe exactly when no feasible future object can tempt a member of the cycle.

**Corollary OA.7** (Support monotonicity)**.** *If $F \subseteq F'$, every cycle safe relative to $F'$ is safe relative to $F$. If the feasible future universe can contain an arbitrary object inserted at the top of any incumbent's ranking, no nontrivial first-round cycle is certifiably safe.*

*Proof.* Both claims follow from Theorem OA.6. Shrinking $F'$ to $F$ only removes upper-contour inequalities. If an arbitrary new object can be placed above a member's current assignment, at least one of those inequalities fails for every nontrivial cycle. □

As the feasible support expands, the set of certified cycles can only shrink; a bound on the future universe is what buys early finality, and without one the frozen-rights conclusion reappears for first-round cycle vesting.

## A.3 A Recursive Certificate Beyond the First Round

First-round cycles survive extension because their members already point to their favorite current objects, whereas a later-round cycle may exist only after earlier objects have disappeared. Certification must therefore proceed in TTC order: a later cycle can be tested only after the preceding rounds have themselves been shown immune to entry.

Run TTC in the current market and let $\mathcal{C}_r$ be the cycles executed in round $r$. Let $x_i$ denote each agent's current TTC assignment. Call a round $F$-satiated if every agent removed in that round prefers $x_i$ to every object in $F$. For a collection $P$ of current TTC cycles, let $g^P$ vest the assignment of every cycle in $P$.

**Proposition OA.8** (Recursive certification)**.** *Suppose rounds $1, \dots, r-1$ are $F$-satiated, and let $P$ be exactly the collection of cycles executed in those rounds. For any nontrivial $C \in \mathcal{C}_r$,*

$$\mathrm{TTC}(E; g^P) = \mathrm{TTC}(E; g^{P \cup \{C\}}) \quad \textit{for every feasible extension } E$$

*if and only if every $i \in C$ prefers $x_i$ to every $z \in F$. Moreover, both allocations equal $\mathrm{TTC}(E; g)$ whenever the condition holds.*

*Proof.* The argument uses cycle stripping twice. Compare $g$ with $g^P$ in an arbitrary extension. After rounds $1, \dots, s-1$ have been removed, a round-$s$ member's assignment is her favorite remaining object from $H$. $F$-satiation places it above every feasible entrant object as well. Under $g$, the round-$s$ members therefore form their current cycles; under $g^P$, they own their assignments and form self-cycles. Either way, removing them leaves the same residual problem. Induction through round $r-1$ gives

$$\mathrm{TTC}(E; g) = \mathrm{TTC}(E; g^P).$$

The cycles in $P$ are also self-cycles under both $g^P$ and $g^{P \cup \{C\}}$, so strip them from that comparison. In the common residual market, $C$ is a first-round cycle. The absorption argument in Theorem OA.6 proves sufficiency when its members rank their assignments above $F$. If $z \succ_i x_i$ for some member, the theorem's one-entrant probe, which uses $i$'s residual right and its predecessor in $C$, gives $z$ to $i$ before $C$ vests and not afterward. This proves necessity. Combining the sufficiency conclusion with the displayed equality proves the final claim. □

The resulting policy is a certified prefix in TTC rounds. Test the cycles in a round separately. Proceed to the next round only if every earlier cycle is future-satiated. Once an earlier cycle fails, later cycles depend on a provisional deletion and the recursion no longer certifies them. Beyond the first round, the exact tests remain Theorem 1 and Corollary OA.2, which cover arbitrary support-specific and history-dependent transitions.

## A.4 Proof Architecture and Constructive Witnesses

The proof turns simultaneous entry into a conditional one-entrant comparison. Given a common tempting probe, order its nonself entrant coordinates and activate them one at a time. Under vested rights, every selected member owns $x_k$; immediately before the first activation that strictly improves any selected member, individual rationality therefore makes every selected member receive exactly $x_k$. If an earlier activation prefix differs across rights, it already separates. Otherwise this preceding profile is a common conditional baseline. Freeze every other entrant report and call the newly activated coordinate pivotal.

At that baseline, the selected members are vested-rights fixed points. Respecting improvement transports the pivotal entrant's effect to a canonical target $x_j$ (Biró et al., 2024). The promoted allocation cycle is a best path from $j$ to that entrant. If it contains another selected $k$, then $k$ is tempted and was removed strictly earlier in the same fixed conditional-baseline execution. Redirecting the same pivotal entrant to $x_k$ lowers this fixed rank. At a terminal path, all other entrants are fixed by the selected ownership rotation. The original-rights conditional baseline gives the selected cohort its current allocation and hence executes its settlement cycle. The same predecessor shortcut as in the one-entrant case then wounds $\pi^{-1}(j)$. Appendix C gives the full argument.

**Example OA.2** (A conditional descent with two active entrants)**.** Let $N = \{1, 2, 3, 4\}$, $H = \{a, b, c, d\}$, original rights be $g = (a, b, c, d)$, and $F = \{u, v\}$. Incumbent preferences are

$$\begin{aligned} 1 &: d \succ u \succ b \succ v \succ a \succ c, \\ 2 &: a \succ c \succ d \succ b \succ u \succ v, \\ 3 &: v \succ d \succ c \succ b \succ a \succ u, \\ 4 &: c \succ d \succ v \succ u \succ a \succ b. \end{aligned}$$

Incumbent-only TTC gives $x = (b, a, d, c)$, first through $C_1 = (3, 4)$ and then through $C_2 = (1, 2)$. Vest both cycles, so $q = (b, a, d, c)$. Write a probe as the pair of targets of the entrants endowed with $u$ and $v$.

At the conditional baseline $\beta = (c, v)$, the first entrant already targets $c$; the second entrant self-cycles. Both rights states give

$$\mathrm{TTC}(\beta; g) = \mathrm{TTC}(\beta; q) = (b, a, d, c, u, v).$$

In a vested-rights execution, member 3 is removed before member 1. Activating the second

entrant toward $b = x_1$ gives the common allocation

$$\mathrm{TTC}((c, b); g) = \mathrm{TTC}((c, b); q) = (d, a, v, c, u, b).$$

Under $q$, its active cycle is $v \to 1 \to 3 \to v$. Thus members 1 and 3 are both tempted, and the latter has the lower baseline removal rank.

Redirect the same entrant to $x_3 = d$. The active cycle shortens to $v \to 3 \to v$, and both rights states give $(b, a, v, c, u, d)$. This path is terminal. The predecessor of 3 in $C_1$ is 4, and $g_3 = c = x_4$. The shortcut that redirects the entrant from $d$ to $c$ gives

$$\mathrm{TTC}((c, c); q) = (b, a, d, c, u, v), \qquad \mathrm{TTC}((c, c); g) = (b, a, v, d, u, c).$$

Member 4 ranks $c$ above $d$, so the shortcut wounds her under original rights and separates. The example displays all three steps of the proof: a genuinely conditional baseline, a strict fixed-rank descent, and a terminal predecessor wound.

**Corollary OA.9** (Constructive conditional witness)**.** *Fix an order of the nonself coordinates of a benchmark probe $a$ that tempts a selected member. A separating probe is found either on the associated activation path or, from the last common neutral prefix $b$, by redirecting the first pivotal entrant to $x_k$ or to $g_k = x_{\pi^{-1}(k)}$ for selected $k$. The construction uses at most $|F| + |V(S)| + 1$ probe evaluations after the benchmark signature is known.*

*Proof.* This is the certificate constructed in the proof of Proposition 3: activation ends at the first selected improvement, canonical redirection visits no selected member twice, and only the final step uses an original-rights target $g_k$. □

**Corollary OA.10** (Singleton three-probe witness)**.** *Suppose $F = \{z\}$. If $S = \{C\}$ and probe $a$ tempts $j \in C$, a separating probe belongs to*

$$\{a, x_j, x_{\pi^{-1}(j)}\}.$$

*If neither $a$ nor $x_j$ separates, the last probe wounds $\pi^{-1}(j)$.*

The characterization in Proposition 3 goes beyond Proposition OA.8: it holds for an arbitrary collection of settlement cycles, whatever the status of earlier TTC rounds. Its proof yields the conditional certificate in Corollary OA.9, and with one future object Corollary OA.10 supplies the sharper three-probe witness.

## A.5 Safety as Zero Worst-Case Welfare Regret

The robust criterion has a Bayesian envelope, but not a fixed-prior welfare ranking. Let $\Omega_F$ be the finite set of strict entrant-preference profiles for the maximal batch $E(F)$, and let $I_F = N \cup E(F)$. For $R \in \Omega_F$ and incumbent rights state $r$, write

$$y^r(R) = \Phi_F^r(R).$$

A normalized cardinal specification $u \in \mathcal{U}_F$ represents every agent's strict preferences by utilities in $[0,1]$, assigning zero and one to her worst and best objects when there is more than one object (and zero to the sole object otherwise). Incumbent utilities are fixed across states; an entrant's utility may depend on her own realized preference type. Let $\Delta^\circ(X)$ denote the strictly positive probability simplex on a finite set $X$. For a full-support prior $\mu \in \Delta^\circ(\Omega_F)$ and strictly positive welfare weights $\lambda \in \Delta^\circ(I_F)$, define

$$W^r(\mu, u, \lambda) = \sum_{R \in \Omega_F} \mu(R) \sum_{i \in I_F} \lambda_i u_i^R(y_i^r(R)) \,.$$

The worst-case Bayesian welfare regret of $q$, relative to retaining $g$, is

$$\overline{\mathrm{Reg}}_F(q) = \sup_{\mu \in \Delta^\circ(\Omega_F),\, u \in \mathcal{U}_F,\, \lambda \in \Delta^\circ(I_F)} \big[W^g(\mu, u, \lambda) - W^q(\mu, u, \lambda)\big]_+, \qquad [d]_+ = \max\{d, 0\}.$$

**Theorem OA.11** (Safety, no harm, and worst-case welfare regret)**.** *For every finite support $F$ and bijective incumbent rights state $q$, the following are equivalent:*

*(i) $q$ is $F$-safe from $g$;*

*(ii) no maximal-batch extension makes any agent worse off under $q$:*

$$y_i^q(R) \succeq_i^R y_i^g(R) \quad \textit{for every } R \in \Omega_F \textit{ and } i \in I_F;$$

*(iii)* $\overline{\mathrm{Reg}}_F(q) = 0$.

*More sharply,*

$$\overline{\mathrm{Reg}}_F(q) = \begin{cases} 0, & q \textit{ is } F\textit{-safe from } g, \\ 1, & q \textit{ is not } F\textit{-safe from } g. \end{cases}$$

*If $q$ is unsafe, there are $R^\star \in \Omega_F$ and $j \in I_F$ such that $y_j^g(R^\star) \succ_j^{R^\star} y_j^q(R^\star)$. Moreover, the strict welfare inequality*

$$W^g(\mu, u, \lambda) > W^q(\mu, u, \lambda)$$

*holds for some full-support prior, normalized compatible utilities, and strictly positive welfare*

*weights.*

*Proof.* Theorem 1 and the maximal-batch reduction imply that $q$ is safe exactly when $y^q(R) = y^g(R)$ for every $R \in \Omega_F$. This immediately implies (ii). Conversely, suppose (ii) holds and fix $R$. The benchmark $y^g(R)$ is a Pareto-efficient TTC allocation. If $y^q(R) \neq y^g(R)$, strict preferences and (ii) make $y^q(R)$ a strict Pareto improvement over $y^g(R)$, a contradiction. Thus (ii) implies equality in every state and hence safety.

Safety makes the two welfare expressions identical for every $(\mu, u, \lambda)$, proving (i) implies (iii). If $q$ is unsafe, the first part gives a state $R^\star$ and agent $j$ who strictly prefers the benchmark allocation. Choose a normalized compatible cardinalization and write

$$\delta = u_j^{R^\star}\left(y_j^g(R^\star)\right) - u_j^{R^\star}\left(y_j^q(R^\star)\right) > 0.$$

Choose an all-positive welfare-weight vector with $\lambda_j \delta > 1 - \lambda_j$. Since every other utility difference is at least $-1$, the benchmark's weighted welfare advantage in state $R^\star$, denoted by $\ell^\star$, is strictly positive. Complete the cardinal specification at all other types arbitrarily. Every state's weighted welfare difference lies in $[-1, 1]$. A full-support prior can put enough probability on $R^\star$ that $\mu(R^\star)\ell^\star > 1 - \mu(R^\star)$. Its expected benchmark advantage is then strictly positive, so $\overline{\mathrm{Reg}}_F(q) > 0$. This proves (iii) implies (i) and the strict-welfare claim.

It remains to establish the stated value under unsafety. Normalized strict representations can make the utility gap between any ordered pair of objects arbitrarily close to one and still fit the finitely many objects above, below, and between them into arbitrarily small neighborhoods of the endpoints. Take such representations with the witness gap tending to one, positive welfare weights with $\lambda_j$ tending to one, and full-support priors with $\mu(R^\star)$ tending to one. The expected benchmark advantage tends to one. Since both weighted welfare levels lie in $[0, 1]$, regret can never exceed one. Its supremum is therefore one. □

The theorem is an envelope statement: an unsafe transition may still be optimal in a specified Bayesian environment, while robust safety rules out every statewise allocation change. With unrestricted ordinal cardinalizations, welfare weights, and nearly concentrated full-support priors, its regret envelope is necessarily binary: it detects unsafety but does not rank unsafe transitions. Once a prior, utilities, and provisionality costs are fixed, the relevant comparison is the expected option value studied in the main article.

The safety test determines when a current trade may become final early. An unsafe trade can still govern current use with the earlier rights retained; Subsection 4.6 of the main article turns that distinction into a clearing rule.

# B Technical Reductions for the Finite-Support Probe Basis

This section states and proves the four technical reductions behind the finite probe basis in Section 4.2 of the main article, which presents their proof architecture and states the resulting safety characterization as a theorem.

## B.1 Peak-report determination

**Lemma OA.12** (Peak-probe determination and rectangular minimality)**.** *Let $I$ be a finite set of agents, $X$ a finite set of personal alternatives, and $\mathcal{R}$ the set of strict orders on $X$. Let $Y$ be a set of complete outcomes, with $y_i \in X$ denoting agent $i$'s assigned alternative at $y \in Y$. Let $f, f' : \mathcal{R}^I \to Y$ be deterministic, strategy-proof, and nonbossy. For every $h \in X$, choose one order $R^h$ that ranks $h$ first and let $B = \{R^h : h \in X\}$. Then*

$$f(R) = f'(R) \quad \text{for every } R \in \mathcal{R}^I \quad \Longleftrightarrow \quad f(R) = f'(R) \quad \text{for every } R \in B^I.$$

*If $|I| = |X| \geq 2$, the basis is cardinality-minimal among rectangular test domains that uniformly determine pairs of one-to-one, strategy-proof, nonbossy assignment rules: each agent's factor must contain a report with every alternative at the top, and hence at least $|X|$ reports.*

*Proof.* Fix agent $i$, opponents' reports $R_{-i}$, and one rule. Let

$$M_i(R_{-i}) = \{f_i(\widetilde{R}_i, R_{-i}) : \widetilde{R}_i \in \mathcal{R}\}$$

be $i$'s attainable menu. Strategy-proofness implies that $f_i(R)$ is the $R_i$-best element of this menu. It also gives the observable menu test

$$h \in M_i(R_{-i}) \quad \Longleftrightarrow \quad f_i(R^h, R_{-i}) = h.$$

The reverse implication is definitional. For the forward implication, if some report attains $h$, a type that ranks $h$ first can attain her unique top alternative; strategy-proofness requires her truthful report $R^h$ to do so. If an arbitrary report and $R^h$ both assign $h$ to $i$, nonbossiness makes their complete outcomes identical.

For $k = 0, \ldots, |I|$, let $\mathcal{D}_k$ contain profiles in which at most $k$ agents report outside $B$. Agreement on $B^I$ is agreement on $\mathcal{D}_0$. Suppose the rules agree on $\mathcal{D}_{k-1}$ and consider a profile in $\mathcal{D}_k$. Replace one nonbasis report $R_i$ in turn by every $R^h$. The resulting profiles lie in $\mathcal{D}_{k-1}$,

so the menu test identifies the same menu for $i$ under both rules. Strategy-proofness assigns the same menu maximum $h$ under $R_i$, and nonbossiness gives

$$f(R_i, R_{-i}) = f(R^h, R_{-i}) = f'(R^h, R_{-i}) = f'(R_i, R_{-i}).$$

Induction proves sufficiency; necessity is immediate.

For minimality, suppose a rectangular test factor for agent $i$ contains no report ranking some object $h$ first. In a one-to-one market, compare two restricted serial rules. Place $i$ first in a fixed serial order. Under the first rule she chooses from all objects; under the second she chooses from $X \setminus \{h\}$. In both rules the remaining agents then choose from all objects left, in the same fixed order. Both rules are strategy-proof and nonbossy. They agree whenever $i$ does not rank $h$ first and differ whenever she does. Thus no such rectangular domain determines every pair of rules. □

**Lemma OA.13** (Elementary TTC probes are complete)**.** *Fix $S \subseteq F$ and two incumbent rights states $q$ and $q'$. Then*

$$\Phi_S^q(R_E) = \Phi_S^{q'}(R_E) \quad \textit{for every strict } R_E$$

*if and only if*

$$\Phi_S^q(P_S^a) = \Phi_S^{q'}(P_S^a) \quad \textit{for every } a : S \to H \cup S.$$

*Proof.* Conditional on incumbent preferences, each rights state induces an assignment rule on the entrants' full strict product domain. TTC is strategy-proof (Roth, 1982) and nonbossy. To see nonbossiness, fix an entrant $e$ and two reports that assign her the same object $h$. Run TTC asynchronously under the two reports. At any common residual market, every cycle that excludes $e$ is available under both reports: only $e$'s outgoing edge can differ. Remove such common cycles until none remains. Cycle-removal order independence implies that $e$ still receives $h$ in both delayed executions, so she points to $h$ under both reports in the resulting common residual market. Every other pointer is also the same; hence the cycle containing $e$ coincides. Removing it leaves the same residual market, and all later assignments coincide. Apply Lemma OA.12, taking $P_z^h$ as the selected top-$h$ report for entrant $e_z$. □

**Lemma OA.14** (The maximal batch subsumes every subbatch)**.** *For any two incumbent rights states $q$ and $q'$,*

$$\Sigma_F(q) = \Sigma_F(q') \quad \Longleftrightarrow \quad \widehat{\Sigma}_F(q) = \widehat{\Sigma}_F(q').$$

*Proof.* Only sufficiency is nontrivial. Fix $S \subseteq F$ and a target map $a : S \to H \cup S$. Extend it

to $\bar{a} : F \to H \cup F$ by setting $\bar{a}(z) = a(z)$ for $z \in S$ and $\bar{a}(w) = w$ for $w \in F \setminus S$. At $P_F^{\bar{a}}$, every added entrant $e_w$ with $w \notin S$ forms a self-cycle with $w$. Remove these cycles first. Active entrants rank the removed objects below their own endowments, and incumbent preferences restrict to $H \cup S$. The residual TTC problem is exactly the subbatch probe $P_S^a$. Equality of the complete maximal-batch outcomes under $q$ and $q'$ therefore implies equality of the residual subbatch outcomes. This holds for every $S$ and $a$. □

## B.2 Entrant-path compression

**Lemma OA.15** (Entrant-path compression)**.** *Fix two incumbent rights states $q$ and $q'$ and a maximal-batch target map $a : F \to H \cup F$. Starting from $z \in F$, follow $a$ whenever the target is another object in $F$. If this path reaches $h \in H$, set $\kappa_a(z) = h$. If it enters a directed cycle contained in $F$, set $\kappa_a(z) = z$. Then $\kappa_a \in \mathcal{A}_F^H$ and*

$$\Phi_F^q(P_F^a) = \Phi_F^{q'}(P_F^a) \quad \Longleftrightarrow \quad \Phi_F^q(P_F^{\kappa_a}) = \Phi_F^{q'}(P_F^{\kappa_a}).$$

*Consequently,*

$$\widehat{\Sigma}_F(q) = \widehat{\Sigma}_F(q') \quad \Longleftrightarrow \quad \widetilde{\Sigma}_F(q) = \widetilde{\Sigma}_F(q').$$

*Proof.* Consider the directed graph on entrants induced by $a$: an entrant points to the owner of her target whenever that target lies in $F$, and a path exits the entrant graph when its target lies in $H$. A component with no old-object exit contains an entrant-only directed cycle. That cycle is a TTC cycle under both rights states. Once it is removed, entrants in the attached in-trees successively lose their targets, point to their own entry objects, and leave in self-cycles. Incoming pointers from incumbents are paths into these cycles and do not alter them. The entire component can therefore be stripped in a way that is independent of incumbent rights. Under $\kappa_a$ the same component is stripped as entrant self-cycles; the residual market is identical.

Now consider all entrant paths ending at some $h \in H$. Until $h$ is removed, these paths form a bundle of gateways to $h$. A directed TTC cycle can enter this bundle through at most one future object $z$. It then follows the unique path from $z$ to $h$ and continues to the incumbent who owns $h$ under the relevant rights state. If a cycle reaches $h$ without entering the bundle, no entrant on these paths joins it. In either case, once $h$ is removed, every unused branch successively falls back to its own object and disappears in self-cycles. Cycle-removal order independence permits these forced cleanup cycles to be removed before the remaining market is continued.

Contract every entrant path $z \rightsquigarrow h$ to the direct gateway $z \to h$. This is exactly the graph induced by $\kappa_a$. Contraction preserves the TTC problem outside the gateway bundle:

it preserves which future object a cycle uses to enter, the old endpoint $h$, and the owner to whom $h$ points.

We now define the expansion map explicitly. Fix a contracted allocation and a component ending at $h$. If no entrant in the component receives $h$, set every entrant in that component to her entry object. Otherwise there is a unique used gateway $z_0$, identified in the contracted allocation by the assignment of $h$ to $e_{z_0}$. Write its unique target path under $a$ as

$$z_0 \to z_1 \to \cdots \to z_\ell \to h.$$

Replace the contracted assignment $e_{z_0} \mapsto h$ by

$$e_{z_0} \mapsto z_1, \quad e_{z_1} \mapsto z_2, \quad \ldots, \quad e_{z_\ell} \mapsto h,$$

with the evident interpretation when $\ell = 0$; every off-path entrant in the component keeps her entry object, and assignments outside the component are unchanged. For a no-exit component, replace the compressed self-allocation by its fixed entrant-cycle allocation under $a$, leaving every attached in-tree entrant at her entry object. Components are disjoint, so these rules define a single rights-independent map $T_a$. Its inverse identifies the used gateway by tracing backward from the entrant who receives $h$ along the entry-object assignments; if no entrant receives $h$, no gateway was used. Thus $T_a$ is injective on the allocations realized as incumbent rights vary, and

$$\Phi^r_F(P^a_F) = T_a(\Phi^r_F(P^{\kappa_a}_F)) \qquad \text{for every incumbent rights state } r.$$

It preserves equality across $q$ and $q'$ in both directions, proving the first equivalence.

Every compressed map belongs to $\mathcal{A}^H_F$, and every $a \in \mathcal{A}^H_F$ satisfies $\kappa_a = a$. Equality on the restricted signature therefore implies equality on every maximal-batch probe; the reverse implication is immediate. □

# C Proof of Finite-Support Benchmark Invariance

Fix a finite future-object set $F$, let $S \subseteq \mathcal{C}$ be an arbitrary collection of nontrivial cycles of the settlement permutation $\pi$, write $V = V(S)$, and let $q = g^S$. Each $z \in F$ has entrant $e_z$ endowed with $z$. Rights assignments extended to all entrants are denoted by $\bar{g}$, $\bar{q}$, and $\bar{r}$. For a probe profile $b$ and entrant $e_z$, write $b[e_z \to h]$ for the elementary probe obtained by changing only that entrant's coordinate to $h \in H \cup \{z\}$.

The proof first selects a conditional baseline at which only one entrant report needs to

change. It then applies fixed-rights promotion tools and a strict descent across the selected fixed points of that baseline.

## C.1 A first-activation baseline

For probe $a \in \mathcal{A}_F^H$, let $A(a) = \{z \in F : a_z \neq z\}$. Choose any order $z_1, \ldots, z_r$ of $A(a)$ and let $a^t$ activate the first $t$ coordinates at their targets in $a$, leaving every later coordinate at self. Thus $a^0$ is the all-self probe and $a^r = a$.

**Lemma OA.16** (First-activation reduction)**.** *Suppose the original-rights benchmark at $a$ tempts a member of $V$. Then some probe on the activation path separates $q$ from $g$, or there are $t \in \{1, \ldots, r\}$ and $j \in V$ such that:*

*(i) $a^{t-1}$ is common under $q$ and $g$, and every $k \in V$ receives $x_k$ there;*

*(ii) $a^t$ is common under $q$ and $g$, and $j$ receives an object strictly preferred to $x_j$ there.*

*Proof.* If $a$ separates, stop. Otherwise its common allocation tempts the same member under $q$. A selected $k$ owns $x_k$ under $q$, so individual rationality places her assignment weakly above $x_k$ at every activation profile. At $a^0$, all entrants self-cycle and Proposition OA.4(ii) gives the incumbent allocation $x$ under both rights states. Let $t$ be the first activation at which some selected member is strictly above $x_k$ under $q$. Strict preferences imply that every selected member receives exactly $x_k$ at $a^{t-1}$. An earlier disagreement is already a separator; otherwise the two displayed profiles are common and have the stated allocations. □

Fix the nonseparating case of the lemma. Call $b = a^{t-1}$ the conditional baseline and $e = e_{z_t}$ the pivotal entrant. Every other entrant report will remain fixed at $b$.

## C.2 Promotion conditional on other entrant reports

**Lemma OA.17** (Conditional target attainment)**.** *Fix rights $r$ and all reports other than $e$'s. Suppose that when $e$ ranks her endowment $z$ first, incumbent $i$ receives her endowment $r_i$. If $e$ ranks $r_i$ first and the resulting TTC allocation gives $i$ an object strictly preferred to $r_i$, then $e$ receives $r_i$.*

*Proof.* While $r_i$ remains, $e$ points to it. Suppose she does not receive it. Because object $r_i$ points to its owner $i$ throughout, the executed cycle that removes $r_i$ contains $i$; it is $i$'s own allocation cycle and gives her the promoted assignment, strictly better than $r_i$ by hypothesis. Every cycle up to and including that one excludes both $e$ and her object. Execute those same cycles at the conditional baseline. The reports, endowments, and available objects relevant to them are unchanged. The singleton $\{e\}$ is available at that baseline. The asynchronous

TTC execution may leave it unchosen and execute the displayed cycles. The cycle containing $i$ would then give her the same strictly better object. TTC order independence contradicts her baseline assignment $r_i$. □

**Lemma OA.18** (Conditional owner targeting)**.** *Under the hypotheses of Lemma OA.17, suppose moving some old object $h$ above $z$ in $e$'s report makes $i$ strictly better than $r_i$. Then the elementary report targeting $r_i$ also makes $i$ strictly better, and $e$ receives $r_i$.*

*Proof.* The case $h = r_i$ is immediate. At the initial improving report, individual rationality gives $e$ either $h$ or $z$. Receipt of $z$ would leave her assignment unchanged from the conditional self-report, so nonbossiness would leave $i$ at $r_i$. Hence $e$ receives $h$.

Insert $r_i$ immediately below $h$ and above $z$. Strategy-proofness gives $e$ the same top object, and nonbossiness preserves the allocation. Now move $r_i$ to first place, leaving every other comparison fixed. This is an improvement of owner $i$'s endowment in the sense of Biró et al. (2024). TTC's respecting-improvement property therefore leaves $i$ strictly above $r_i$. Lemma OA.17 implies that $e$ receives $r_i$ at this modified profile.

At the elementary report $r_i \succ_e z \succ_e \cdots$, the entrant can obtain her top object $r_i$ by reporting the modified ranking. Strategy-proofness therefore gives her $r_i$ at the elementary report, and nonbossiness leaves the rest of the modified allocation unchanged. Incumbent $i$ remains strictly above $r_i$. □

## C.3 Active promotion paths

At a baseline where $e$ ranks her endowment $z$ first, use the following *entrant-last execution*. As long as a TTC cycle excluding $e$ is available, execute one such cycle. Execute the always-available singleton $\{e\}$ only when no other cycle is available, and then complete TTC arbitrarily. Construct $G_{\mathrm{CP}}$ from this one-cycle-at-a-time execution as in Biró et al. (2024): record every cycle edge and, when a cycle is removed, every current top-choice edge from a surviving agent into that cycle. A *best path* from $i$ to $e$ follows at each vertex the most-preferred recorded edge among those that can still reach $e$.

The entrant-last convention matters. If the available singleton $\{e\}$ is executed prematurely, a later active promotion need not be represented by a path in that particular $G_{\mathrm{CP}}$. The convention fixes one admissible baseline execution, independent of the owner subsequently targeted, and is used for every canonical redirection below.

**Lemma OA.19** (Cycle precedence)**.** *If $G_{\mathrm{CP}}$ contains a path from $u$ to $v$, then either $u$ and $v$ lie in the same executed cycle or $v$ is removed before $u$.*

*Proof.* This is Fact 2 of Biró et al. (2024, p. 1946). It also follows directly from the construction. A cycle edge joins agents removed in the same cycle. A predecessor edge is recorded from a surviving agent into the cycle currently being removed, so its head is removed strictly before its tail. Traverse a directed path. If every edge is a cycle edge, all its vertices belong to one cycle. Otherwise, the first predecessor edge makes removal time fall strictly, and subsequent edges can only preserve or further reduce it. Hence the terminal vertex is removed before the initial one. □

**Lemma OA.20** (Active promotion path)**.** *Fix rights $r$ and all other reports. At a baseline profile, $e$ ranks her endowment $z$ first. Move incumbent $i$'s endowment $r_i$ above $z$ in $e$'s ranking, leaving all other relative comparisons unchanged. If the new TTC allocation gives $e$ object $r_i$, its allocation cycle containing $e$ consists of a best path from $i$ to $e$ in the baseline $G_{\mathrm{CP}}$ and the closing edge $e \to i$.*

*Proof.* First we show that the entrant-last baseline graph contains a path from $i$ to $e$. Suppose instead that the baseline execution removes the cycle containing $i$ before it removes $\{e\}$. Replay, at the promoted profile, the baseline cycles through and including that cycle. None contains $e$, and the only changed pointer is $e$'s, so every replayed cycle remains available. The replay removes object $r_i$ without $e$. TTC cycle-order independence then contradicts the hypothesis that the promoted allocation assigns $r_i$ to $e$. Hence $i$ is still present when the entrant-last execution removes $\{e\}$.

Immediately before that removal, no cycle excluding $e$ is available. The current top-pointer path starting at every surviving agent therefore ends at the singleton $\{e\}$. Follow the path from $i$. Its last edge into $e$ is recorded when $\{e\}$ is removed. After that removal, the predecessor of $e$ is eventually removed in a cycle. Its own predecessor still points to it, so the preceding edge is recorded either as a cycle edge or as a predecessor edge. Iterating backward records the entire path from $i$ to $e$ in $G_{\mathrm{CP}}$. Since $e$ always points to $z$ at the baseline, her only recorded outgoing edge is her self-cycle edge. Thus $i$ and $e$ are in distinct baseline cycles and $i$ is a noncycle predecessor of $e$.

The conclusion is now exactly Case III in the proof of Theorem 1 of Biró et al. (2024, p. 1947). Their improved owner $i$ is incumbent $i$ here, their unique preference-changing agent $j$ is entrant $e$, object $i$ is $r_i$, and $j$'s baseline allotment is $z$. Case III states that the promoted trading cycle is the baseline best path from $i$ to $e$ closed by the new edge $e \to i$, as required. □

## C.4 Descent across conditional-baseline fixed points

Fix once and for all the entrant-last one-cycle-at-a-time execution at $b$ under $q$. Every selected $k \in V$ owns and receives $x_k$, so $k$ is a singleton cycle. Let $\rho_b(k)$ be its removal step. For a common allocation $w$, define

$$\sigma_q(i) = \bar{q}^{-1}(w_i), \qquad \sigma_g(i) = \bar{g}^{-1}(w_i),$$

and let $\Pi_S$ equal $\pi$ on $V$ and the identity elsewhere.

**Lemma OA.21** (Conditional canonical descent)**.** *Suppose $b[e \to x_j]$, $j \in V$, has a common allocation $w$ under $q$ and $g$, and $w_j \succ_j x_j$. Let $L_j$ be the $q$-allocation cycle containing $e$. Then $e$ receives $x_j$, and $L_j$ consists of a best path from $j$ to $e$ in the fixed conditional-baseline graph $G^q_{\mathrm{CP}}(b)$ and the closing edge $e \to j$. Moreover, every $k \in (L_j \cap V) \setminus \{j\}$ satisfies*

$$\rho_b(k) < \rho_b(j) \qquad \textit{and} \qquad w_k \succ_k x_k.$$

*Proof.* Lemmas OA.17 and OA.20 give the first two statements. A selected $k \neq j$ on $L_j$ lies on the best path from $j$ to $e$, so $G^q_{\mathrm{CP}}(b)$ contains a path from $j$ to $k$. The two agents are distinct singleton cycles at $b$. The path-order property therefore gives $\rho_b(k) < \rho_b(j)$. Because $L_j$ is nontrivial, $\sigma_q(k) \neq k$ and hence $w_k \neq q_k = x_k$. Individual rationality under $q$ makes the inequality strict in the favorable direction. □

**Lemma OA.22** (Conditional terminal shortcut)**.** *Under the hypotheses of Lemma OA.21, suppose*

$$L_j \cap V = \{j\}.$$

*Let $p = \pi^{-1}(j)$. At probe $b[e \to g_j] = b[e \to x_p]$, the original-rights benchmark wounds $p$. Hence that probe separates $q$ from $g$.*

*Proof.* Since $q_v = g_{\Pi_S(v)}$, the common allocation satisfies

$$\sigma_g = \Pi_S \circ \sigma_q. \tag{2}$$

Write the entrant cycle under $q$ as

$$L_j : \quad e \to j \to u_1 \to \cdots \to u_s \to e.$$

None of $u_1, \ldots, u_s, e$ is selected. Equation (2) therefore preserves the path $j \to u_1 \to \cdots \to u_s \to e$ and changes the closing edge to $e \to \pi(j)$. Thus $j$ and $e$ lie in the same $g$-allocation cycle $B$.

The conditional baseline $b$ is common and gives every selected member her current object. Under original rights, the complete settlement cycle $C$ containing $j$ is therefore an executed TTC cycle: its allocation-owner map is $k \mapsto g^{-1}(x_k) = \pi(k)$ for every $k \in C$, so $C$ is a closed allocation cycle with those edges. The owner of $x_j$ is $\pi(j)$, so Lemma OA.20 represents $B$ by a best path from $\pi(j)$ to $e$ in the original-rights baseline graph $G^g_{\mathrm{CP}}(b)$ and the edge $e \to \pi(j)$. Since $j \in B$, this path reaches $j$. Before doing so it cannot leave $C$ and later re-enter: the path-order property would require an outside cycle both to precede and to follow $C$. Along its segment in $C$, the path's only recorded within-cycle edge is the executed cycle edge. It consequently contains

$$\pi(j) \to \pi^2(j) \to \cdots \to p \to j.$$

Combining this segment with the preserved suffix of $L_j$ gives

$$B: \quad e \to \pi(j) \to \pi^2(j) \to \cdots \to p \to j \to u_1 \to \cdots \to u_s \to e. \tag{3}$$

Choose an asynchronous original-rights execution in which $B$ is a trading cycle at $b[e \to x_j]$. At $b[e \to g_j]$, execute the same cycles preceding $B$. They exclude $e$, so changing only $e$'s outgoing pointer does not destroy them. Immediately before $B$ would execute, the entrant points to $g_j$, owned by $j$; at the canonical probe she points to $x_j$, owned by $\pi(j)$. The cycle in (3) is shortened to

$$e \to j \to u_1 \to \cdots \to u_s \to e.$$

It removes $g_j = x_p$ and leaves $p$ in the market.

At the corresponding canonical stage, $p$ pointed to $x_p$ along the edge $p \to j$ in $B$. Every object that $p$ prefers to $x_p$ had therefore already been removed. The shortcut probe removes those same objects and then removes $x_p$ without $p$, so $p$ eventually receives an object strictly worse than $x_p$. Under $q$, $p$ owns $x_p$; individual rationality gives her an object weakly above it. The probe separates. □

Other active entrants may lie among $u_1, \ldots, u_s$ or in cycles preceding $B$. They are fixed by $\Pi_S$, and their reports do not change between the canonical and shortcut probes. The terminal argument keeps them in the market and lets their trades intertwine with $B$.

*Proof of Proposition 3.* Suppose benchmark neutrality fails. If a selected member is wounded, Proposition OA.4(i) supplies a separator. Otherwise some probe $a$ tempts a selected member. Apply Lemma OA.16. A separating activation prefix ends the argument. In the remaining case, let $b$ be the common neutral conditional baseline and $e$ the pivotal entrant. At the next activation some selected $j$ is strictly improved. Lemma OA.18 transports the improvement to $b[e \to x_j]$. If that probe separates, stop; otherwise it satisfies Lemma OA.21.

If $L_j \cap V = \{j\}$, Lemma OA.22 supplies the separating probe $b[e \to g_j] = b[e \to x_{\pi^{-1}(j)}]$. Otherwise choose any $k \in (L_j \cap V) \setminus \{j\}$. Lemma OA.21 makes $k$ tempted and gives $\rho_b(k) < \rho_b(j)$. Redirect the same pivotal entrant to $x_k$ using Lemma OA.18. A separating canonical probe ends the argument; a common one repeats the step at $k$. The conditional baseline and every background report remain fixed, so every continuation uses the same rank $\rho_b$. It falls strictly, and the process terminates after at most $|V|$ canonical redirections, followed, if necessary, by one terminal shortcut. Thus the post-activation stage evaluates at most $|V| + 1$ probes, and every failure of benchmark neutrality yields a separator.

Conversely, benchmark neutrality excludes temptation at every probe. Part (iii) of Proposition OA.4 gives safety. The condition is memberwise and uses only the original-rights benchmarks, so it holds for a union exactly when it holds for every constituent cycle. This proves factorization. □

*Proof of Corollary OA.10.* Let $F = \{z\}$ and $S = \{C\}$. The activation baseline is the all-self probe. In its vested-rights execution, remove external predecessor cycles and then choose the fixed point $\{j\}$ before the other fixed members of $C$. The path-order property prevents the best path from $j$ to $e$ from containing a member of $C \setminus \{j\}$. Hence $L_j \cap V = \{j\}$, and Lemma OA.22 applies immediately. □

## C.5 Proof of state-by-state retention

*Proof of the state-by-state parts of Proposition 3.* Part (iv) implies (iii). If (iii) holds, it holds at every maximal-batch compressed probe, which is part (ii); the equivalence of (i) and (ii) proved above then gives (i).

Suppose (i) holds. That equivalence assigns $x_k$ to every selected member at every compressed maximal-batch probe under original rights. Fix an arbitrary supported subbatch and entrant profile. Pad each absent entrant with a self-cycle; this leaves the original extension as the residual TTC problem. Let $y$ be the resulting maximal-batch allocation under $g$.

For each entrant $e_z$, replace her report by the elementary top report that ranks her realized assignment $y_{e_z}$ first and her entry object $z$ second when they differ. The realized object belongs to the entrant's attainable menu. Strategy-proofness therefore assigns it again when it is made her unique top object, and nonbossiness leaves the complete allocation unchanged. Repeating entrant by entrant gives a maximal elementary target map $a : F \to H \cup F$ with allocation $y$.

The map may contain entrant-to-entrant target paths. Entrant-path compression contracts every path ending at an old object and replaces every entrant-only component by self-targets. Its rights-independent expansion changes only entrant coordinates, so the compressed probe

has the same incumbent allocation as $P_F^a$. The compressed map belongs to $\mathcal{A}_F^H$; hence every selected $k$ receives $x_k$. Removing the padded self-cycles proves (iii) for the original extension. Finally, safety equates the complete allocations under $g$ and $g^S$, so (iii) implies (iv). □

# D The Two-Protector Reduction

*Proof of the hardness part of Theorem 4.* Reduce minimum vertex cover. Let $D = (P, L)$ be a finite simple undirected graph with at least one edge. Create a protector cycle $P_v$ for every vertex $v \in P$ and an investment cycle $T_\ell$ for every edge $\ell \in L$. Every cycle has two agents and one target. There are

$$K = |P| + |L|$$

equiprobable states: one protector-cost state per vertex and one investment hazard state per edge.

For edge $\ell = \{u, v\}$, the hazard state contains the interaction ring

$$T_\ell \longrightarrow P_u \longrightarrow P_v \longrightarrow T_\ell.$$

Its incoming demands at $P_u$ and $P_v$ are nonaligned, and its incoming demand at $T_\ell$ is aligned. Its exposed set is therefore

$$E_\ell = \{P_u, P_v\}.$$

The target investment at $T_\ell$ is insured exactly when at least one endpoint cycle vests. Utilities in Appendix E make the ring's policy-dependent allocation gain zero.

In vertex state $v$, attach a nonaligned self-loop entrant to $P_v$ and give the live trade allocation gain $K$. Since the state has probability $1/K$, leaving $P_v$ provisional yields one unit of expected allocation value. Vesting it forgoes that value. Every policy-dependent allocation gain is thus nonnegative and direct waiting costs are zero.

Only targets of the $T_\ell$ cycles can invest. Give each the quadratic indirect value $v(\rho) = \theta\rho^2/2$. If neither endpoint of $\ell$ vests, retention is $1 - 1/K$; if either vests, retention is one. Fix coverage reward $M = 2$ and choose

$$\theta = \frac{2M}{1 - (1 - 1/K)^2}.$$

The investment gain from covering $\ell$ is then exactly $M$.

Give the target of every $T_\ell$ a costless separate investment licence and leave $T_\ell$ provisional. Because its incoming demand is aligned, replacing this licence by title would leave the state

allocation and the target's retention unchanged; title at $T_\ell$ is therefore irrelevant to the reduction. Every feasible policy consequently has a welfare-equivalent policy with all $T_\ell$ cycles provisional. Let $S \subseteq P$ be the protector cycles that vest. After dropping constants, expected integrated welfare is

$$M\big|\{\ell \in L : S \cap \ell \neq \varnothing\}\big| - |S|.$$

If an edge is uncovered, adding either endpoint raises the first term by at least two and costs one. Every optimizer is a vertex cover. Conditional on covering all edges, the first term is constant, so every optimizer is a minimum vertex cover.

The construction has rational probabilities and utilities with polynomial encoding length. The explicit TTC implementation below uses canonical two-agent cycles and one gateway per cycle. A policy-independent aligned auxiliary trade in every state adds a strict allocation gain without changing the displayed objective. Hence terminal allocation welfare strictly exceeds the no-entry benchmark in every state and under every policy. □

# E TTC and Utility Implementations

For every protector and investment cycle, write its members as $(q^0, q^1)$. Agent $q^0$ is the target. She owns $h_{q^0}$ and currently receives $h_{q^1}$; agent $q^1$ owns $h_{q^1}$ and currently receives $h_{q^0}$. Normalize original-object utilities to zero. Fix $G > \theta$. Give every target value $G$ for her current object and every nontarget value $3G$ for hers. Rank all unmentioned objects below the original object in a strict order.

In edge state $\ell = \{u, v\}$, orient the interaction ring as $T_\ell \to P_u \to P_v \to T_\ell$. Each target values her state-specific new object at $2G$. Every entrant values her demanded object at $G$ and her own new object at zero. The demand entering $T_\ell$ is its target's current object and is aligned. The demands entering $P_u$ and $P_v$ are their targets' original objects and are nonaligned.

When the ring is live, all three targets receive their new objects, all three entrants receive their demands, and the nontarget in $T_\ell$ retains her current object. The protector nontargets receive original objects. Live welfare is

$$3(2G) + 3G + 3G = 12G,$$

equal to current surplus $3(G + 3G) = 12G$. Vesting either protector makes its nontarget self-cycle with the demanded original object, breaks the ring, and restores every incumbent's current assignment. Vesting $T_\ell$ alone preserves the ring because its incoming demand is

aligned.

In vertex state $v$, the target of $P_v$ values a new object at $2G$. Her entrant makes the nonaligned demand for the target's original object and values it at $2G + K$. If $P_v$ waits, trade welfare is $4G + K$, compared with current surplus $4G$; if $P_v$ vests, the trade is blocked. The allocation gain is therefore $K$.

For strict state-policy gains, add one fixed auxiliary two-agent cycle. In every state, its target values a state-specific entrant object at $2G$; the entrant makes an aligned demand for the target's current object and values it at $G$; the nontarget keeps a current object worth $3G$. The trade produces $6G$, compared with current surplus $4G$, independently of vesting.

For each investment target, effort lies in $[0, \theta]$, gross adaptation return is $e$ upon retention, and cost is $e^2/(2\theta)$. The value of $\theta$ in Appendix D makes the insured-versus-uninsured investment difference exactly two. Since $G > \theta$, even maximal adaptation return cannot reverse

$$z_q \succ h_{q^1} \succ h_{q^0}.$$

The implementation is rank preserving. Taking $G = \theta + K + 1$ gives rational utilities of polynomial bit length; small distinct negative values for unmentioned objects complete strict rankings.

## E.1 The information environment

For a singleton information mode, use one two-agent investment cycle $Q$ and one two-agent protector cycle $P_j$. The hazard state contains the interaction ring

$$Q \longrightarrow P_j \longrightarrow Q.$$

The incoming edge at $Q$ is aligned and the incoming edge at $P_j$ is nonaligned. Leave $Q$ provisional and give its target a costless separate investment licence. Replacing the licence by title at $Q$ would not block the ring or change the target's retention. Vesting $P_j$ blocks the ring. For a general tail $E_j$, put every tail member on the same ring as a nonaligned destination and retain $Q$ as its aligned, licensed head.

The ring's allocation gain $B_j$ is calibrated through entrant utilities without moving a TTC pointer. The same quadratic investment technology and the restriction $G > \theta$ preserve rankings. This implements the hyperedges, retention losses, and allocation gains in Proposition 6. The additive $\gamma_e$'s in that proposition are conditional reduced-form opportunity costs: whenever protector $e$ is selected, $\gamma_e$ is incurred after either an ex ante or a mode-contingent title decision. In particular, the public mode signal does not itself erase the

cost. Auxiliary nonaligned cost states can generate such terms in environments with a suitably separated conditional state space, but the proposition treats those auxiliary costs as reduced-form primitives.

### E.2 Executable verification

The replication package contains the support-safety checks and the bundled-authority endpoint tests. The hybrid-authority script separately checks the fixed- and endogenous-licence title margins, singleton-tail supermodularity with arbitrary licence costs, the loss of the general head–protector complementarity under free licences, exact minimum-vertex-cover recovery with every investment head licensed for free, and the aligned-head TTC implementation. The remaining scripts check the information formulas, asymmetric counterexample, and support refinements. These exercises verify the displayed constructions, finite-domain identities, and computational reductions; the analytic proofs stand on their own.

## F The Repeated Full-Vesting Benchmark

Repeated full-vesting TTC runs TTC from current rights at every clearing and vests the resulting allocation fully. This section states and proves the properties summarized in Section 3 of the main article.

**Proposition OA.23** (Repeated full-vesting TTC)**.** *Under repeated full-vesting TTC, for every finite history:*

*(i) every clearing is individually rational relative to current rights and Pareto efficient;*

*(ii) each incumbent weakly improves along the realized path; and*

*(iii) the terminal allocation is individually rational relative to original rights and Pareto efficient in the final market.*

*The rule therefore satisfies spot Pareto efficiency and latest-implemented-assignment protection. It can still violate both integration-order and clearing-refinement neutrality, and its terminal allocation can lie outside the core defined by original rights.*

*Proof.* Static TTC gives individual rationality and Pareto efficiency at each clearing. Because full vesting makes the current allocation the next outside-option vector, individual rationality generates the monotone path in (ii); transitivity then gives individual rationality relative to original rights. The final TTC run gives terminal Pareto efficiency.

In Example 1 of the main article, arrival of agent 2 before agent 3 produces $(b, a, c)$; arrival of agent 3 before agent 2 produces $(c, b, a)$. Integration-order neutrality therefore fails. Along

the second path, deferring the clearing after agent 3's arrival produces the one-shot TTC allocation $(b, a, c)$, so clearing-refinement neutrality also fails. At $(c, b, a)$, coalition $\{1, 2\}$ can use its original objects $\{a, b\}$ and reassign them as $(b, a)$, strictly improving both members. The repeated outcome is therefore blocked relative to original rights. □

# G Periodic Settlement TTC

This section states the implementation result summarized in Subsection 4.6 of the main article. Fix exogenous settlement dates that partition time into epochs. In the first epoch, reference rights are the initial endowments. At the start of every later epoch, each incumbent's reference right is her allocation at the preceding settlement. Entry makes the record append-only within an epoch: each new agent adds her entry object without changing an existing reference right.

At a provisional clearing, the institution runs TTC for the active market from the current reference record and may implement the allocation for current use. The provisional allocation is not written back into that record. At settlement, TTC is run from the same accumulated record on all agents then active. The settlement allocation becomes the reference-right state of the next epoch. Call this rule *periodic-settlement TTC*.

**Proposition OA.24** (Clearing without continuous vesting)**.** *Fix a settlement calendar and exogenous entry. Suppose no incumbent or old object exits and incumbent preferences extend to new objects without reversing comparisons among previously available objects. Periodic-settlement TTC satisfies:*

*(i) conditional on the final extension in an epoch, its settlement allocation is clearing-refinement neutral and integration-order neutral;*

*(ii) every provisional and settlement allocation is individually rational relative to the then-current reference rights and Pareto efficient in the active market; and*

*(iii) it satisfies settled-assignment protection.*

*Proof.* Within an epoch, provisional computations do not change rights, and entry objects are appended independently of arrival order. For a fixed final extension, settlement therefore faces the same agents, objects, preferences, and rights after every provisional-clearing schedule and entrant order. This proves part (i). Every run is ordinary TTC from the reference rights then present, so static individual rationality and Pareto efficiency give part (ii).

For part (iii), let $x^\tau$ be the allocation at settlement $\tau$. It becomes the reference-right state for the next epoch. Individual rationality gives every agent active at $\tau$ an object weakly preferred to $x_i^\tau$ at each provisional clearing and at settlement $\tau + 1$. The latter allocation

then becomes the next reference state. Iteration, together with preservation of comparisons among previously available objects, protects $x_i^\tau$ at every later clearing. An entrant's first allocation is likewise weakly preferred to her entry object because that object is her initial reference right. □

Two histories may implement different provisional allocations and still reach the same settlement allocation. The rule obtains path neutrality by leaving provisional assignments open to reversal until settlement. At a fixed settlement market, it also inherits TTC's static incentive property: truthful reporting is dominant for an agent who values only her settlement object, conditional on the agents, objects, and reference rights available at that settlement.

# References


Biró, P., Klijn, F., Klimentova, X., and Viana, A. (2024). Shapley–Scarf housing markets: Respecting improvement, integer programming, and kidney exchange. *Mathematics of Operations Research*, 49(3):1938–1972.

Roth, A. E. (1982). Incentive compatibility in a market with indivisible goods. *Economics Letters*, 9(2):127–132.

Roth, A. E. and Postlewaite, A. (1977). Weak versus strong domination in a market with indivisible goods. *Journal of Mathematical Economics*, 4(2):131–137.