

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

Dalin Sheng*

August 2026

Abstract

Dynamic assignment rules must choose whether a current allocation becomes the endowment of a later market. We study this rights transition in finite Shapley–Scarf economies with entry. A rights state is support-safe when it produces the same terminal allocation under top trading cycles as original rights for every supported extension and entrant report. Safety is characterized by a finite basis of elementary entrant probes. For transitions that vest current trading cycles, safety factorizes cycle by cycle and yields a unique maximal safe settlement. We then separate exchange title from licences that authorize relationship-specific investment. On a canonical stochastic domain, future trades induce a directed title-insurance hypergraph: each head is an exposed investment and its tail contains the titles that block reassignment. Singleton tails make the planner’s reduced objective supermodular for every licence cost even when allocation gains have either sign; with free licences and quadratic investment, one minimum cut finds an optimum. Tails of size two already generate substitution and NP-hardness. Public information creates value by identifying the title that protects a licensed investment.

Keywords: dynamic matching; top trading cycles; investment licences; property rights; title insurance.

JEL classification: C78, D47, D61.

1 Introduction

An allocation reached today can become the endowment of tomorrow’s market. Suppose top trading cycles (TTC) assigns an incumbent a new object and further endowed agents may

*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).

enter. The institution can implement the assignment for current use while retaining the old rights record, or it can vest the assignment and use it as the endowment of the next TTC. We ask which rights transitions leave every supported future allocation unchanged and how the institution should choose among transitions that do affect future trade.

The welfare question becomes sharper when current use, investment authority, and exchange rights are separate instruments. A user can occupy an object and receive a licence to adapt it while another rights state governs the next market. Relationship-specific investment pays only if the user retains the adapted object. Exchange title can raise this retention probability by blocking a later trading path. The title that supplies protection may sit on another TTC cycle, since its economic role is to break the path that would reassign the investor. Vesting therefore supplies network insurance and simultaneously removes an exchange option.

The distinction appears literally in the decant rules of Welwyn Hatfield Borough Council. A tenant temporarily moved from home a to home b signs a licence for b and retains the tenancy at a (Welwyn Hatfield Borough Council, 2025a, secs. 13.1, 13.13, and 19.1). A permanent decant requires panel approval and the signing of a new tenancy at b before the move is finalized (Welwyn Hatfield Borough Council, 2025a, secs. 19.2 and 19.4). Occupation, authorization, and the exchangeable tenancy record move on distinct dates. Section 7 gives the full institutional chronology and relates the legal record to mutual exchange.

Our model is a finite Shapley–Scarf economy with entry. An intermediate TTC allocation determines current use. Before the future extension is realized, the institution selects a bijective rights state for the next TTC and may issue separate investment licences. Investment preserves ordinal rankings, so the rights state determines its return through the probability that the investor retains the current assignment. The safety analysis allows arbitrary strict incumbent preferences and arbitrary entrant reports on a finite support. The welfare analysis then specializes the future exchange network to canonical current cycles and stochastic single-gateway entry. This specialization converts retention risk into an exact network object while retaining the rights transitions characterized in the general model.

Two benchmark results locate the problem. First, immediate protection of every implemented assignment creates an entitlement ratchet. With three agents, spot efficiency and protection force different terminal allocations under two orders of entry (Proposition 1). Second, if the institution requires a rights transition to be invisible to every possible one-agent extension, the rights state must remain fixed (Proposition 2). A separating entrant offers a new object for an incumbent object whose recorded owner changed. These results explain why useful early finality requires restrictions on future entry.

Theorem 1 gives the first main result. Fix a finite support F of possible entrant objects

and allow every subset to arrive with arbitrary strict entrant preferences. A candidate rights state is F -safe when terminal TTC agrees with TTC from original rights in every such extension. Safety is equivalent to agreement on one finite family of elementary probes: all supported entrants arrive together, and each entrant either keeps her object or ranks one incumbent object above it. The proof combines a menu argument for strategy-proof and nonbossy assignment rules, padding of missing entrants by self-cycles, and compression of entrant paths to incumbent targets. The basis contains $(n + 1)^m$ profiles for n incumbent objects and $m = |F|$. It is polynomial in market size for fixed m , while its exponential dependence on support size reflects the fact that joint entrant batches can reveal a rights change that every singleton batch misses.

For the policy class that vests unions of current TTC cycles, Proposition 3 converts the two-rights comparison into a one-sided test. A collection is safe exactly when every selected member keeps her current assignment in every elementary benchmark computed from original rights. The necessity argument is the substantive step: a probe that merely tempts a selected member need not separate the two rights states, so the proof follows a finite promotion path until it finds a wounding probe. Theorem 2 then shows that safety factorizes across current cycles. One set of benchmarks identifies a unique maximal safe settlement, and this set expands monotonically when information removes possible entrant objects. Earlier settlements remain safe after every support refinement.

Safety also settles the case where title spends no exchange option. On a safe cycle, the relevant users retain their assignments in every supported state under either rights record. Full title and a separate licence induce the same terminal allocation and first-best investment. Relative to optimized licensing, title either saves a licence cost that would be paid or restores the investment surplus when that cost would deter licensing. A free licence makes the two instruments welfare-equivalent (Proposition 4).

The unsafe region requires an exact account of future exchange paths. Theorem 3 supplies it on the stochastic single-gateway domain. Entrant demands concatenate into directed interaction rings. Each ring that exposes investment q creates a hyperedge $E \rightrightarrows q$: the head is the exposed investment, and the tail E contains the exchange titles that block the ring. A licence activates the head without changing its tail. The resulting *title-insurance hypergraph* determines all retention probabilities and the planner’s reduced objective, including the case of free licences. Proposition 5 then gives the instrument margin:

$$\text{saved licence cost} + \text{network insurance gain} \geq \text{destroyed exchange option}.$$

The network term includes investments on other cycles whose licences become valuable when

title raises retention.

Theorem 4 identifies a sharp structural boundary. When every nonempty hyperedge tail is a singleton, retention and allocation welfare are modular in titles, even when allocation gains have either sign. Convex investment value then makes the reduced objective supermodular for every vector of licence costs, so titles are complements and an optimum is computable in polynomial time. With quadratic investment and either free or prohibitively costly licences, one minimum cut suffices. A tail with two protectors reverses the cross-difference: the protectors are substitutes. Tails of size two also make welfare maximization NP-hard with canonical two-agent cycles, free licences, nonnegative allocation gains, and rank-preserving quadratic investment. Redundant protection changes a closure problem into a covering problem over exchange rights.

Section 6 uses the hypergraph to value information available before title is chosen. For one freely licensed investment, Proposition 6 separates the value of adapting effort from the value of adapting the protector set. In the symmetric fragmented case, ignorance produces an all-or-none title decision, while disclosure targets the single realized protector (Corollary 1). In the intermediate region, disclosure changes title policy even though effort-only information has zero value. For a general hypergraph in the full-insurance region, the legal value of information is the gap between the cost of an ex ante minimum hitting set and the expected cost of hitting the realized tail (Corollary 2).

The timing implication follows from the same distinction between records. Safe title can move immediately. A licence can authorize adaptation while an unsafe exchange record remains provisional. Periodic settlement provides one implementation: allocations may change within an epoch, and the rights state moves only at the settlement boundary. The Online Appendix states the formal properties of this implementation.

The closest matching literature treats existing rights as an input to the allocation rule. TTC starts from an ownership record (Shapley and Scarf, 1974; Ma, 1994); house allocation with existing tenants protects current assignments (Abdulkadiroğlu and Sönmez, 1999); and a recent reallocation mechanism uses an initial matching as its individual-rationality benchmark and weakly improves a distributional objective (Hafalir et al., 2025). Work on changing populations and market integration likewise studies allocation conditional on the rights inherited by the enlarged market (Ehlers et al., 2002; Kumar et al., 2022). Our object is the transition rule that creates those inherited rights. The probe characterization identifies all transitions that are observationally equivalent on a given entry support, and the cycle result constructs the largest safe settlement.

Several dynamic matching papers make current assignments relevant to later markets; Doval (2025) provides a recent survey. Kurino (2014) compares spot mechanisms with

and without property-rights transfer in an overlapping-generations housing market; under transfer, a newcomer’s current assignment becomes her next-period endowment. [Matsui and Murakami \(2022\)](#) analyze an initial deferred-acceptance allocation followed by retrade. [Combe et al. \(2025\)](#) characterize efficient and incentive-compatible spot mechanisms in large dynamic markets without money, and [Nicolò et al. \(2025\)](#) develop dynamic core notions and an intertemporal TTC. Dynamic applications use TTC or priorities to govern reassignment ([Huitfeldt et al., 2026](#)), including future chain externalities created by kidney-exchange vouchers ([Ashlagi et al., 2026](#)). We hold the terminal rule at TTC and endogenize the rights state passed into it. This isolates the allocation effect of a rights transition and produces an exact cyclewise safety test.

The investment block connects this transition problem to the design of property rights. Ownership affects noncontractible investment in incomplete-contract models ([Grossman and Hart, 1986](#); [Hart and Moore, 1990](#)), while empirical work links property rights to investment incentives ([Besley, 1995](#)). [Bloch and Houy \(2012\)](#) assign a durable object to successive agents under temporary property rights; [Weyl and Zhang \(2022\)](#) design depreciating licences to balance investment and reallocation; and [Dworczak and Muir \(2025\)](#) design menus of property rights that determine outside options in later trade. Sunk investment can also shape the matching that follows ([Nöldeke and Samuelson, 2024](#)). Here investment authority is available separately from exchange title, including at zero cost. Title acts through the future TTC network and can protect an investment on a different cycle. The hypergraph representation identifies exactly where this protection comes from.

The network interpretation complements work that treats endowments as exclusion rights ([Balbuzanov and Kotowski, 2019](#)) and traces exclusion through production and trading networks ([Balbuzanov and Kotowski, 2024](#)). Related market-design applications study ownership transfers to public-housing tenants and rights that govern trade externalities ([Andersson et al., 2016](#); [Ferguson and Milgrom, 2026](#)). Our representation maps future TTC rings into protecting-title sets and exposes a rank boundary: unique protection produces complementarity and polynomial optimization, while two substitute protectors generate a covering problem.

The information results connect to disclosure in dynamic allocation ([Ashlagi et al., 2025](#)) and to the standard Blackwell value of information ([Blackwell, 1953](#)). The signal here identifies the legal instrument relevant to a licensed investment. The welfare decomposition separates effort adjustment from title adjustment, and the hypergraph recourse formula prices the second component through a hitting-set gap.

Section 2 defines rights transitions and the two path comparisons. Section 3 gives the entitlement-ratchet benchmark. Section 4 characterizes support-safe transitions, constructs

maximal cyclewise settlement, and introduces licences and investment. Section 5 derives the title-insurance hypergraph and its tractability boundary. Section 6 studies contingent title after public information. Section 7 gives the institutional interpretation and records the domain of each result.

2 Rights Transitions and Timing

Current use and future exchange rights are separate records. An allocation gives the first; a rights state gives the agent to whom each object points when TTC next runs. Clearing can change use without changing rights. A transition then decides whether and how the allocation vests, so the same pattern of use can lead to different trades after entry.

2.1 Objects, allocations, and rights

A finite housing market contains agents N and indivisible objects H , with $|N| = |H|$. Agent i has a strict preference $\succ_i$ over H , with $\succeq_i$ denoting its weak extension. Strictness makes TTC single-valued and removes the need for a tie-breaking rule.

A *rights state* is a bijection $g : N \rightarrow H$; g_i is the object controlled by agent i . We call i its owner under g : object g_i points to i in the TTC graph. Ownership in this sense is the mechanism's exclusion-right record. An *allocation* is another bijection $x : N \rightarrow H$; x_i is the object used by agent i . The model permits $x \neq g$: an agent can occupy one object and control another.

An allocation x is individually rational relative to g if $x_i \succeq_i g_i$ for every i . It is Pareto efficient if there is no feasible y such that $y_i \succeq_i x_i$ for every i and $y_j \succ_j x_j$ for some j .

At a TTC clearing, every remaining agent points to her favorite remaining object and every remaining object points to its owner under g . All directed cycles execute and are removed, and the procedure repeats. We write $\text{TTC}(N, H, g, \succ)$ for the resulting allocation. On the strict housing-market domain, TTC is individually rational, Pareto efficient, and strategy-proof, and it selects the strict-core allocation (Shapley and Scarf, 1974; Roth and Postlewaite, 1977; Roth, 1982; Ma, 1994).

2.2 Entry, clearing, and rights transitions

The initial market is $(N_0, H_0, g^0, \succ)$. A future extension is *balanced*: it adds endowed agents and the same number of distinct objects, and no incumbent or old object leaves. Each entrant initially controls one new object and has a strict preference over the enlarged object set. Incumbent preferences also extend to the enlarged set, preserving every comparison between

old objects; entry and preference profiles are exogenous. This no-reversal condition keeps the rights effect of entry separate from preference reversal, which would change allocations even under fixed rights. The unrestricted benchmark uses one entrant, denoted by 0, with entry object z ; the finite-support and stochastic analyses also permit simultaneous batches. For any realized extension, *original rights* combine g^0 for the initial incumbents with the object brought by each entrant.

Fix a current market $(N, H, g, \succ)$. Let $\mathfrak{h}$ record public information available before the next extension, including earlier arrivals, implemented and settled allocations, and any public forecast of future entry. A clearing rule selects an allocation x , and a *rights transition* T selects the bijection used at the next clearing:

$$\tilde{g} = T(\mathfrak{h}, N, H, g, \succ, x).$$

Full vesting sets $\tilde{g} = x$; frozen rights set $\tilde{g} = g$. More generally, T may condition on the observed history and select any bijection, whether or not it follows the current allocation. Implementation and vesting are therefore separate operations: x may govern current use, with $\tilde{g} = g$ retained as the reference-right state.

The clearing-refinement and investment comparisons use the following timing.

- (1) In market $(N, H, g, \succ)$, an intermediate clearing selects x , which may be implemented immediately for current use.
- (2) Public information available at the rights date is observed. The institution selects exchange rights $\tilde{g}$ and a set of separate investment licences before the future extension.
- (3) Users authorized either by title or by a separate licence choose investment.
- (4) New endowed agents and their objects enter, and incumbent preferences extend to the new objects.
- (5) The enlarged market is cleared by TTC from $\tilde{g}$ together with the entrants' initial rights; investment returns are realized.

Definition 1 (Nonanticipatory transition). A rights transition is nonanticipatory if, conditional on the public information available before the future extension, $\tilde{g}$ and every separate licence are fixed before the realized extension is observed.

Pre-entry information may include current preferences and allocations, the support or distribution of future entry, and incumbents' rankings over named future objects when these have been elicited in advance. In the unrestricted benchmark, comparisons with an

object outside the known current market are part of the future extension. Nonanticipation, the restriction in Proposition 2, excludes a designer who observes the realized entrants, endowments, and preferences and then decides whether the previous trade vested.

The robust comparisons in Sections 3–4 are ordinal and concern the terminal allocation. Cardinal utilities enter only when the analysis values allocation gains and investment. Keeping the layers separate lets the safety results identify the states in which title has no allocation opportunity cost before the planner prices unsafe finality.

2.3 Protection and path neutrality

A dynamic rule selects an implemented allocation after each history. We use two path comparisons. Integration-order neutrality holds the final market fixed and varies the order in which entrants join. Clearing-refinement neutrality holds the extension fixed and varies whether an intermediate clearing changed rights. The first comparison applies to any dynamic rule, whereas the second fixes the terminal rule at TTC. Both require a precise account of what an earlier assignment protects.

Definition 2 (Latest-implemented-assignment protection). A dynamic rule satisfies latest-implemented-assignment protection if (i) an entrant’s first implemented assignment is weakly preferred to her entry object, and (ii) after every implemented allocation x_t , each agent i active at date t receives $x_s(i) \succeq_i x_t(i)$ at every later date $s > t$.

This condition makes implementation itself the source of a new outside option. Because assignments can only improve under part (ii), the latest one is the binding floor. It is the strong protection premise used in the ratchet result. A weaker *settled-assignment protection* condition applies part (ii) only to an allocation explicitly declared settled. Periodic settlement satisfies the weaker condition, though it violates the stronger one whenever a provisional assignment is reversed.

Definition 3 (Spot efficiency). A dynamic rule is spot Pareto efficient if every implemented allocation is Pareto efficient for the agents and objects then active.

Definition 4 (Integration-order neutrality). Fix an initial market and a balanced final extension: the entrants, their entry objects, and the strict preferences of every final agent over the final object set. A dynamic rule is integration-order neutral if its terminal allocation is unchanged when only the entrants’ arrival order varies.

Definition 5 (Clearing-refinement neutrality). Fix a current history, a future extension, and a final TTC clearing. Compare an *early* history in which an intermediate clearing triggers a

rights transition $g \mapsto \tilde{g}$ with a *deferred* history in which that intermediate clearing is omitted and rights remain g . Both histories have the same final agents, objects, and preferences. The transition is clearing-refinement neutral for the extension if final TTC selects the same allocation in the two histories. It is *universally one-entry neutral* if this holds after every current history and for every extension by one endowed agent.

Section 3 uses the first comparison to isolate the entitlement ratchet without committing to a dynamic allocation rule. Section 4 then fixes TTC and uses the second comparison to ask which intermediate rights changes are invisible to future entry.

3 A Benchmark: Immediate Protection

Vesting every efficient assignment upon implementation protects current improvements, but it also lets early arrival determine the outside options protected against later entry. We first isolate this force for any deterministic dynamic rule, and then examine repeated full-vesting TTC, the canonical rule that equates current use with the rights state of the next market.

Example 1 (A three-agent entry race). Agent 1 initially controls a . Agents 2 and 3 can enter with b and c . Preferences are

Agent	Initial right	Preference
1	a	$b \succ_1 c \succ_1 a$
2	b	$a \succ_2 b \succ_2 c$
3	c	$a \succ_3 b \succ_3 c$

If agent 2 arrives first, spot efficiency in the two-agent market uniquely selects (b, a) for agents (1, 2). Both agents then hold their top objects. Latest-implemented protection keeps these assignments fixed, and agent 3 retains her entry object, giving terminal allocation (b, a, c) .

If agent 3 arrives first, spot efficiency selects (c, a) for agents (1, 3). When agent 2 arrives, protection keeps a with agent 3, requires agent 1 to receive at least c , and gives the newcomer at least her entry object b . Feasibility now forces terminal allocation (c, b, a) .

Spot efficiency pins down the first trade on each path. Protection then turns that trade into a path-specific floor: the two orders reach the same final market yet leave no common terminal allocation consistent with both floors. The impossibility therefore rests on the conditions alone.

Proposition 1 (Immediate-protection ratchet). *On the unrestricted strict-preference domain, no deterministic dynamic rule satisfies spot Pareto efficiency, latest-implemented-assignment protection, and integration-order neutrality for every balanced extension from one initial incumbent. Three agents suffice and are minimal.*

Proof. The same agents, objects, original rights, and preferences in Example 1 lead to different forced terminal allocations under the two entrant orders. Hence integration-order neutrality fails.

For minimality, with at most two agents there is at most one entrant. At entry, choose, using a fixed tie-break, a Pareto-maximal allocation subject to both agents receiving at least their original rights. The constraint set is nonempty, and any Pareto improvement of a feasible allocation remains in it; hence the choice is spot efficient. It protects the incumbent’s singleton assignment and the entrant’s entry object. With no subsequent assignment or alternative entrant order, all three conditions hold. Thus at least three agents are necessary. $\square$

Both forced terminal allocations are Pareto efficient. The conflict is between protection and integration-order neutrality: the first efficient assignment contracts the set of later trades that every incumbent will accept. Protection is the irreversible step; implementation alone creates no such floor.

Reversibility removes the displayed conflict. After the first entrant arrives in Example 1, implement the efficient two-agent swap provisionally and leave original rights unchanged. Once all three agents are present, run TTC from those rights. Both arrival paths then produce (b, a, c) . This repairs the example under settled-assignment protection, and Section 4.6 develops the repair into a general rule: frequent implementation need not create the ratchet when provisional use can be reversed.

Repeated full-vesting TTC is the natural benchmark that refuses this reversibility. At every clearing it computes the static TTC allocation and makes that allocation the next rights state. The benchmark shows which static properties survive this updating and which path-independence properties do not.

The benchmark keeps every static guarantee. Each clearing is individually rational relative to current rights and Pareto efficient; because full vesting makes the current allocation the next outside-option vector, each incumbent weakly improves along the realized path; and 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. What fails is path neutrality. In Example 1, the two arrival orders produce the terminal allocations (b, a, c) and (c, b, a) , so integration-order neutrality fails;

deferring the second path’s intermediate clearing restores (b, a, c) , so clearing-refinement neutrality fails as well; and at (c, b, a) the coalition $\{1, 2\}$ can reassign its original objects $\{a, b\}$ as (b, a) to the strict benefit of both members, so the terminal allocation lies outside the core defined by original rights. Online Appendix F states these properties as a proposition and proves them.

The positive and negative conclusions have the same source. Full vesting resets the individual-rationality benchmark after every clearing. The reset makes assignments improve along the realized path; it also makes the benchmark path dependent. The terminal allocation can therefore be Pareto efficient and individually rational relative to every preceding assignment, yet blocked relative to original rights. Static TTC guarantees do not aggregate into dynamic path neutrality.

The ratchet arises from immediate assignment protection. The next section holds the future extension fixed and allows an arbitrary rights transition. It asks whether any change $g \mapsto \tilde{g} \neq g$ can be invisible to future entry.

4 Support-Safe Rights Transitions

From now on we hold the final extension fixed and compare an early clearing that changes rights with a deferred clearing that leaves the same current rights in place. The analysis proceeds from the worst case to the institutional case. With unrestricted entry, neutrality permits no change in rights (Section 4.1). With a known finite support, arbitrary rights states are equivalent exactly when they agree on a finite family of elementary entry probes (Section 4.2). For pure settlement, that is, vesting collections of current TTC cycles, the probe benchmarks computed under original rights characterize safety exactly, cycle by cycle, and deliver the maximal safe settlement together with its information ratchet (Sections 4.3 and 4.4).

4.1 Robust benchmark: unrestricted entry

The strongest robustness requirement is neutrality to every one-agent extension. The intermediate transition may be arbitrary: it may record any rights state, including states that cut across the current TTC cycles. The following construction shows that a suitable entrant can reveal any pairwise change in the legal record.

One-entry separation. Fix two distinct incumbent rights states g and $\tilde{g}$. Some old object h has different owners: let i own it under g and $j \neq i$ own it under $\tilde{g}$. Add one entrant 0

endowed with a new object z . Insert z at the top of both i 's and j 's rankings and at the bottom of every other incumbent's ranking, preserving every comparison among old objects. Let entrant 0 rank

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

The remaining old objects may appear in any strict order. Under g , the first-round TTC graph contains

$$i \longrightarrow z \longrightarrow 0 \longrightarrow h \longrightarrow i,$$

so i receives z . Under $\tilde{g}$, it contains the analogous cycle through j , so j receives z . Other agents can point into these cycles without changing the assignments of their members.

The construction is the option-transfer mechanism at its starkest. The entrant's object is an offer for the selected old object h , and the owner of h is the counterparty able to accept it. A difference in the legal record therefore becomes a difference in the terminal allocation, though every incumbent comparison between old objects is untouched. One entrant is also the tight worst-case bound: with no entry, distinct states need not be distinguishable. If $a \succ_1 b$ and $b \succ_2 a$, TTC assigns (a, b) both from rights (a, b) and from the swapped rights state.

Proposition 2 (Universal safety characterizes frozen rights). *A nonanticipatory rights transition is universally one-entry clearing-refinement neutral if and only if it leaves the rights state unchanged at every intermediate clearing.*

Proof. If every transition leaves rights unchanged, inserting an intermediate clearing leaves the agents, objects, final preferences, and rights in the final TTC problem unchanged. Neutrality follows.

Conversely, suppose that after some observed history the transition changes rights from g to $\tilde{g} \neq g$. Condition on that complete history. Nonanticipation fixes $\tilde{g}$ before the next entrant is observed. The one-entry construction above supplies an extension for which final TTC from g and final TTC from $\tilde{g}$ differ. The deferred history reaches the final clearing with g ; the early history reaches it with $\tilde{g}$. Universal neutrality is contradicted. Thus no intermediate transition can change rights. $\square$

The universal quantifier does the work. The institution may still implement a current allocation and let the assigned user occupy the object; what the proposition constrains is the legal record, because when the next trading opportunity is unrestricted, every nontrivial transfer of rights can matter.

The conclusion extends beyond TTC taken as a primitive.

Axiomatic implication. Suppose every final clearing must be individually rational, Pareto efficient, and strategy-proof on the unrestricted strict housing-market domain. For each fixed rights state, Ma (1994) identifies TTC as the unique rule satisfying these three axioms. Proposition 2 therefore implies that, among nonanticipatory transitions, universal one-entry neutrality is equivalent to freezing rights at every intermediate clearing.

Ma’s result selects the allocation conditional on rights; the one-entry construction disciplines the transition between rights states. The robust conclusion relies on admitting the separating object. A known support, however, may exclude that object and leave room for early finality. The rest of the paper works out how much room.

4.2 An exact probe basis under finite support

Fix a current market $M = (N, H, g, \succ)$ and a finite set F of objects that may enter before settlement. Each $z \in F$, if it arrives, comes with a new owner. Incumbents’ strict preferences over $H \cup F$ are known when the rights decision is made. The extension domain contains every subset of F and every strict preference profile for the corresponding entrants.

Testing a candidate rights state against every extension and every strict entrant report is finite but factorially large. This subsection reduces the test to a sparse finite basis.

For $S \subseteq F$, let $E(S) = \{e_z : z \in S\}$ contain one entrant endowed with each $z \in S$. For an incumbent rights state $q : N \rightarrow H$ and entrant preference profile R_E , write $\Phi_S^q(R_E)$ for the complete allocation selected by TTC in the market with agents $N \cup E(S)$, objects $H \cup S$, incumbent rights q , and entrant rights $e_z \mapsto z$.

Fix an arbitrary reference order of $H \cup F$ and use its restriction in each market $H \cup S$. For entrant e_z and target $h \in H \cup S$, define the *elementary probe* P_z^h by placing z first when $h = z$, and by placing h first and z second when $h \neq z$; all remaining objects follow below z in the reference order. Thus a probe entrant either keeps her entry object or is willing to exchange it for exactly one target. The target may be another entrant’s object, and several entrants may target the same object. For a target map $a : S \rightarrow H \cup S$, let $P_S^a = (P_z^{a(z)})_{z \in S}$. The reference order below z is immaterial: as long as e_z remains active, her owned object z also remains, so her TTC pointer never reaches an object ranked below z .

Definition 6 (Support safety and probe signature). A rights state q is *F-safe from* g if

$$\Phi_S^q(R_E) = \Phi_S^g(R_E)$$

for every $S \subseteq F$ and every strict entrant-preference profile R_E . Its *F-probe signature* is the indexed family

$$\Sigma_F(q) = (\Phi_S^q(P_S^a))_{S \subseteq F, a: S \rightarrow H \cup S}.$$

Its *maximal-batch signature* is

$$\widehat{\Sigma}_F(q) = (\Phi_F^q(P_F^a))_{a:F \rightarrow H \cup F}.$$

Let

$$\mathcal{A}_F^H = \{a : F \rightarrow H \cup F : a(z) \in H \cup \{z\} \text{ for every } z \in F\}$$

and define the *incumbent-target signature*

$$\widetilde{\Sigma}_F(q) = (\Phi_F^q(P_F^a))_{a \in \mathcal{A}_F^H}.$$

These signatures record the full allocation, including incumbent and entrant assignments.

The reduction behind the characterization applies to any pair of strategy-proof, nonbossy assignment rules; the TTC structure only makes the determining reports economically sparse, with each entrant finding at most one trade acceptable relative to keeping her endowment. Here strategy-proofness concerns the agent's assigned alternative, and nonbossiness means that a unilateral report change that leaves her assignment unchanged leaves the complete allocation unchanged.

Proof architecture. The reduction has three steps. First, a menu argument shows that the peak reports, one top- h report per alternative h , determine any deterministic, strategy-proof, nonbossy assignment rule. Strategy-proofness identifies the best element of an agent's report-independent option set; nonbossiness allows off-basis reports to be replaced one at a time without changing the complete outcome. When the number of reporting agents equals the number of alternatives, this rectangular basis is cardinality-minimal.

Second, conditional on incumbent preferences, TTC under any fixed rights state is strategy-proof and nonbossy as a rule of entrant reports. The menu argument therefore reduces arbitrary entrant preferences to the elementary probes. Missing entrants can then be added as self-cycles, so equality on the maximal batch implies equality on every subbatch. Third, entrant-to-entrant target paths carry no additional information about incumbent rights. A chain ending at an old object is a single gateway to that object, whereas a component that never reaches an old object trades only among entrants. Contracting each chain to its old endpoint, and expanding it by a rights-independent map, reduces the maximal-batch signature to $\widetilde{\Sigma}_F$. Online Appendix B states and proves the four technical reductions.

Theorem 1 (General finite-support safety characterization). *For every finite incumbent Shapley–Scarf economy, every finite future support F , and every bijective incumbent rights*

state q ,

$$\begin{aligned} q \text{ is } F\text{-safe from } g &\iff \Sigma_F(q) = \Sigma_F(g) \iff \widehat{\Sigma}_F(q) = \widehat{\Sigma}_F(g) \\ &\iff \widetilde{\Sigma}_F(q) = \widetilde{\Sigma}_F(g). \end{aligned}$$

Consequently, all support-safe rights transitions are given exactly by

$$\mathcal{Q}_F(g) = \{q \in \text{Bij}(N, H) : \widetilde{\Sigma}_F(q) = \widetilde{\Sigma}_F(g)\}.$$

If a transition is unsafe, it can therefore be separated by the market containing every supported entrant, with each entrant either keeping her entry object or willing to exchange it for one incumbent object.

Proof. Safety implies equality at every probe profile. Conversely, equality of $\widetilde{\Sigma}_F$ expands to equality on every maximal-batch target map. Padding absent entrants with self-cycles gives equality on every subbatch, and peak-report determination extends this equality from elementary probes to every strict entrant-preference profile. This is F -safety. The same reductions in reverse give all displayed equivalences. $\square$

Remark 1 (Joint batches are indispensable). Safety against each singleton support $\{z\}$, $z \in F$, does not imply F -safety. Online Appendix A constructs a three-agent market and a rights change that survive every one-entrant probe yet become separated when two supported entrants arrive together: the second entrant removes the object through which the first entrant's trade would otherwise route, and the rerouted trade reaches the incumbent whose recorded rights differ. The signatures in Theorem 1 therefore quantify over joint target maps. The maximal-batch reduction shows that the full batch alone carries all of the information.

The characterization also has an exact no-harm interpretation. A bijective rights state q is F -safe from g if and only if, for every $S \subseteq F$, every strict entrant profile R_E , and every agent i of the extended market,

$$\Phi_S^q(R_E)_i \succeq_i \Phi_S^g(R_E)_i,$$

where $\succeq_i$ is i 's preference in that extension. Safety gives the inequalities with equality. Conversely, if they hold and the allocations differ, strict preferences make some agent strictly better off under q , with no one worse off. This contradicts Pareto efficiency of the original-rights TTC allocation. Thus, state by state, a transition that never harms anyone cannot change anyone's assignment.

Online Appendix A sharpens the equivalence into an envelope statement: the normalized worst-case Bayesian welfare regret of a transition is zero if it is safe and one if it is not. The robust criterion therefore detects unsafety but does not rank unsafe transitions. Ranking them requires the priors and cardinal utilities used in Section 5.

4.3 Safe vesting of current cycles

The general signature permits arbitrary bijective changes in rights. We now specialize to the institutional decision that motivates the paper: which parts of the current TTC allocation may vest. Let $x = \text{TTC}(N, H, g, \succ)$ and define the settlement permutation

$$\pi(i) = g^{-1}(x_i), \quad \text{so that} \quad x_i = g_{\pi(i)}.$$

Let $\mathcal{C}$ be the nontrivial cycles of π . For $S \subseteq \mathcal{C}$, write $V(S)$ for the union of its members and let g^S assign x_i to each $i \in V(S)$ and retain g_i elsewhere. Because $V(S)$ is a union of complete π -cycles, g^S is bijective. The incumbent-only TTC allocation remains x : in an original TTC execution, each selected trading cycle, once reached, can be replaced by the self-cycles of its members, leaving the same residual market.

For an elementary probe $a \in \mathcal{A}_F^H$, let $y^a = \Phi_F^g(P_F^a)$ denote the original-rights benchmark.

Proposition 3 (Safe vesting under finite support). *For every finite F and every collection $S \subseteq \mathcal{C}$, the following are equivalent:*

- (i) g^S is F -safe from g ;
- (ii) $y_k^a = x_k$ for every $a \in \mathcal{A}_F^H$ and $k \in V(S)$;
- (iii) every selected member $k \in V(S)$ receives x_k under original rights at every supported extension and entrant report;
- (iv) every selected member receives x_k under both g and g^S at every supported extension and entrant report.

Thus a pure settlement is safe exactly when every selected member keeps her current assignment in every original-rights elementary benchmark, or equivalently in every supported state.

Proof idea. The difficult step is the equivalence of (i) and (ii). A wounding benchmark already separates the two rights states. If a benchmark only tempts a selected member, the proof redirects the pivotal entrant along a sequence of selected fixed points until either a probe separates or the entrant demands the assigned object of a tempted member's cycle predecessor, which wounds that predecessor. Parts (ii)–(iv) then follow by peak-report determination, path compression, and nonbossiness: an arbitrary entrant profile can be

replaced by elementary reports without changing the complete allocation. Online Appendix C gives the full construction.

The state-by-state equivalence settles what vesting adds. On a safe cycle, waiting already provides de facto certainty that the selected users retain their current objects. Vesting cannot raise investment by reducing retention risk there. A separate licence can supply the same authority. Safe title can therefore save a licence cost, although with costless licences it has no additional investment value.

The proposition turns equality between two allocation rules into a one-sided test on original-rights counterfactuals. The institution never simulates a candidate settlement; it runs the elementary probes once, from the rights it already has, and reads off which cycles may vest. A cycle can vest only when no supported entry configuration changes the object that any of its members carries into the next market.

For a selected member k and an elementary probe a , call the benchmark *neutral*, *tempting*, or *wounding* at k according as $y_k^a = x_k$, $y_k^a \succ_k x_k$, or $x_k \succ_k y_k^a$. A tempting benchmark shows a supported entrant luring k away from her cycle trade; a wounding benchmark shows entry destroying that trade to k 's detriment. The next example illustrates the test, the option-transfer mechanism behind it, and the reason its necessity direction is delicate.

Example 2 (Vesting as an option transfer). Four incumbents 1, 2, 3, 4 own a, b, c, d , respectively. Two future objects z and w are conceivable. Preferences over $H \cup \{z, w\}$ are

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

Incumbent TTC executes two first-round cycles: C_1 swaps a and b between agents 1 and 2, and C_2 swaps c and d between agents 3 and 4, so $x = (b, a, d, c)$.

Suppose first that only z can arrive: $F = \{z\}$. An elementary probe sends one entrant, endowed with z , who either keeps z or offers it for one incumbent object. The five original-rights benchmarks are:

probe target	benchmark y^a for (1, 2, 3, 4; entrant)	classification
a	$(b, a, d, c; z)$	neutral at every member
b	$(b, a, d, c; z)$	neutral at every member
c	$(b, a, z, d; c)$	tempts 3, wounds 4
d	$(b, a, z, c; d)$	tempts 3
z	$(b, a, d, c; z)$	neutral at every member

Members 1 and 2 keep their current assignments in every benchmark, so Proposition 3 certifies C_1 : whether agent 1 owns a or b at the next clearing is invisible to every supported extension. Member 3 does not, so C_2 cannot vest, and no collection containing C_2 can.

The two failing probes teach different lessons. The d -probe tempts 3 without separating. Under original rights the entrant's offer routes through a three-way trade in which 3 takes z , the entrant takes d , and 4 takes c ; after vesting C_2 it routes through a bilateral trade between 3 and the entrant, and 4 keeps her vested object c directly. Yet the allocation is $(b, a, z, c; d)$ either way. A tempting benchmark warns of trouble; temptation alone does not establish unsafety. The c -probe separates. Under original rights, 3 sells her endowment c to the entrant, and her cycle predecessor 4, who was slated to receive c , is stranded with d : the benchmark wounds 4. After vesting, c belongs to 4, who ranks it first and refuses to sell; the entrant keeps z and 3 keeps d . The allocations $(b, a, z, d; c)$ and $(b, a, d, c; z)$ differ, so the probe demanding the predecessor's assigned object is a separator. This is the proof of the proposition in miniature: a temptation somewhere becomes detectable at the tempted member's predecessor.

Now enlarge the support to $\{z, w\}$. Member 1 ranks w above her assignment b , and the probe offering w for a (with the z -entrant keeping z) separates C_1 : under original rights, agent 1 sells a and takes w ; after vesting C_1 , object a belongs to agent 2, who ranks it first, and the trade dies. The option to sell a to a future w -supplier is worth exercising to agent 1 but worthless to agent 2; vesting moves the option between them, and a supported entrant can tell. The maximal safe settlements are

$$S^*(\{z, w\}) = \emptyset, \quad S^*(\{z\}) = \{C_1\}, \quad S^*(\emptyset) = \{C_1, C_2\}.$$

If public announcements first rule out w and later rule out z , cycle C_1 vests at the first announcement and C_2 at the second. Nothing settled ever reopens.

Benchmark screening. Two pointwise implications organize the proof. If a probe wounds a selected member, vested-rights individual rationality separates the two allocations. If a probe tempts no selected member, replacing every vested-right arc in a putative blocking cycle by its original-right path would block the original-rights TTC allocation; hence the probe cannot separate. A collection that is never tempted is therefore safe. Online Appendix A states and proves these implications for arbitrary collections.

Why the benchmark test is exact. Sufficiency follows from benchmark screening: if every selected coordinate equals x_k at every benchmark, no probe tempts a selected member. Necessity is the substantive direction. A wounding benchmark already separates, so the

remaining case is a probe that tempts a selected member while wounding none, and, as the d -probe in Example 2 shows, such a probe need not separate on its own. The proof locates a separator conditionally. Activate the probe’s nonself entrant coordinates one at a time. At the first activation that strictly improves a selected member under vested rights, either an earlier activation prefix already separates the rights states or the preceding profile is a common conditional baseline at which vested-rights individual rationality pins every selected member to her current assignment. Holding all other entrant reports fixed turns the problem into a one-entrant comparison. Promoting the pivotal entrant’s target through the selected fixed points of that baseline traces a promotion path; each failed separator hands the path to a selected member removed strictly earlier in a fixed baseline execution, so the removal rank strictly decreases; the baseline and all background reports stay put. Finiteness forces either a separator or a terminal probe in which the pivotal entrant demands the assigned object of a tempted member’s cycle predecessor: this shortcut wounds the predecessor, exactly as in the example, and separates by benchmark screening. Online Appendix C develops the promotion lemmas, the descent invariant, and the complete proof; Online Appendix A works a two-entrant example in which the descent takes every one of these steps.

4.4 The maximal settlement and the information ratchet

Because the test in Proposition 3 is memberwise and its benchmarks do not depend on the candidate settlement, safety of a collection is equivalent to safety of its parts.

Theorem 2 (Maximal safe settlement and information ratchet). *For every finite F :*

(i) *for every $S \subseteq \mathcal{C}$,*

$$g^S \text{ is } F\text{-safe} \iff g^{\{C\}} \text{ is } F\text{-safe for every } C \in S.$$

(ii) *there is a unique inclusion-maximal safe collection $S^*(F)$. It contains exactly the current cycles whose members keep their current assignments at every original-rights benchmark. If $n = |H|$, $m = |F|$, and TTC is implemented in quadratic time, it is found in $O((n+1)^m(n+m)^2)$ time without enumerating the $2^{|C|}$ cycle collections;*

(iii) *if $F' \subseteq F$, then*

$$S^*(F) \subseteq S^*(F').$$

More generally, for nested supports $F_0 \supseteq F_1 \supseteq \dots \supseteq F_T$, put $q_t = g^{S^(F_t)}$. Then q_t is F_t -safe from q_{t-1} for every $t \geq 1$. Information can therefore add safely vestable cycles without reopening an earlier settlement.*

Proof. The condition in Proposition 3 is memberwise. It holds for a union of disjoint cycles exactly when it holds for each cycle. Taking the union of all cycles that satisfy it gives the unique maximal safe collection. The elementary basis contains $(n + 1)^m$ probes; each original-rights benchmark is computed once and its member coordinates classify all current cycles simultaneously.

If $F' \subseteq F$, every extension supported on F' is admissible under F after entrants in $F \setminus F'$ are omitted, or equivalently made self-cycles. Every F -safe cycle is therefore F' -safe, and maximality gives the inclusion. For nested supports, q_{t-1} and q_t are both F_t -safe from g , so their terminal TTC allocations agree at every F_t extension. $\square$

Factorization makes the characterization usable. A priori, safety of a settlement could depend on which other cycles vest alongside it; the theorem says it cannot, so one family of benchmarks classifies every cycle and every union at once. The $(n + 1)^m$ term is the size of the joint probe basis and is polynomial in the market size n for any fixed support size m , but the exponential dependence on m reflects Remark 1: joint batches separate rights states that no singleton batch does. For arbitrary bijective rights changes the characterization still applies state by state, though the exact procedure here enumerates candidate bijections. Online Appendix A records the resulting bound together with polynomial verification of any single proposed transition for fixed m .

Part (iii) is the entitlement ratchet in its constructive form. Section 3 showed that protecting assignments as they are implemented ratchets away future trades. Here information, not implementation, runs the ratchet: finality expands as uncertainty about future demand resolves, and finality granted this way never has to be clawed back. Example 2 displays the resulting settlement path.

The exact test also has a closed-form specialization that requires no probe computation at all. If C is a nontrivial first-round current TTC cycle, then C is safe relative to F if and only if

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

This condition makes the cycle absorbing under either rights state: every member already points to her current assignment, entry or no entry, so both histories remove the same agents and objects and leave the same residual market. If it fails at (i, z) , an entrant endowed with z who demands i 's original object gives z to i under original rights, whereas after vesting i does not receive z : it either goes to i 's cycle predecessor or remains with the entrant. In Example 2, cycle C_1 satisfies the condition for $F = \{z\}$ and fails it for $F = \{z, w\}$ through the pair $(1, w)$. Online Appendix A extends the condition to a recursive certificate for later TTC rounds along a certified prefix, records the blocking-cycle and common-core representations

of general safe transitions, and gives a constructive separating witness whose length is linear in the number of supported objects and the number of selected members.

4.5 Safe title and separate licences

The safety results are ordinal. We now add relationship-specific investment to price the rights transitions they characterize. Exchange rights and investment authority are separate policy variables.

Assumption 1 (Title and separate investment authority). *If the selected rights state satisfies $q_i = x_i$, user i holds title to her current assignment. Title makes x_i the object she supplies to the next TTC and automatically authorizes adaptation. If $q_i \neq x_i$, the institution may issue a separate investment licence at resource or verification cost $d_i \geq 0$. The licence authorizes the same adaptation without changing any ownership pointer in the next TTC.*

At $d_i = 0$, authority can be fully separated from exchange title. Positive costs make separation an endogenous choice.

Fix a distribution μ over future extensions. Conditional on a selected rights state q , let $Y^q(\omega)$ be the terminal TTC allocation in state ω . Let $L \subseteq \{i : q_i \neq x_i\}$ be the separately licensed users. A user is authorized when $q_i = x_i$ or $i \in L$, and chooses effort $e_i \in E_i$ after the policy decision and before ω . The adaptation yields gross return $b_i(e_i)$ if $Y_i^q(\omega) = x_i$ and zero otherwise, and costs $k_i(e_i)$. The zero action is feasible and satisfies $b_i(0) = k_i(0) = 0$. Define

$$\rho_i(q; \mu) = \mathbb{P}_\mu\{Y_i^q(\omega) = x_i\}, \quad v_i(\rho) = \max_{e \in E_i} \{\rho b_i(e) - k_i(e)\}.$$

If E_i is compact and b_i, k_i are continuous with $b_i \geq 0$, then v_i exists, is nonnegative, nondecreasing, and convex: it is the pointwise maximum of affine, nondecreasing functions of ρ . Under $b_i(e) = e$, $k_i(e) = e^2/(2\theta_i)$, and $e \in [0, \theta_i]$,

$$e_i^*(\rho) = \theta_i \rho, \quad v_i(\rho) = \frac{\theta_i}{2} \rho^2.$$

When $d_i \geq v_i(1)$, a separate licence is never strictly optimal; omitting licences from the menu gives the bundled-title benchmark. Let $A_\mu(q)$ denote expected allocation welfare, with policy-independent constants suppressed. Integrated welfare is

$$\Psi_\mu(q, L) = A_\mu(q) + \sum_{i: q_i = x_i \text{ or } i \in L} v_i(\rho_i(q; \mu)) - \sum_{i \in L} d_i. \quad (1)$$

For fixed exchange rights, licence decisions separate. An unvested user is licensed exactly

when $v_i(\rho_i(q; \mu)) \geq d_i$, with both actions optimal at equality. Writing $[z]_+ = \max\{z, 0\}$, the reduced exchange-rights objective is

$$\tilde{\Psi}_\mu(q) = A_\mu(q) + \sum_i [\mathbf{1}\{q_i = x_i\}v_i(\rho_i(q; \mu)) + \mathbf{1}\{q_i \neq x_i\}[v_i(\rho_i(q; \mu)) - d_i]_+]. \quad (2)$$

Assumption 2 (Rank preservation). *Investment does not change any strict object ranking used by TTC.*

A sufficient condition is that the largest feasible adaptation return is smaller than the smallest cardinal utility gap between objects in the fixed ranking. Rank preservation lets the ordinal allocation results determine retention, while investment supplies a cardinal welfare consequence. Private and social adaptation returns coincide, and investment has no salvage value to its maker after reassignment. Section 7 records the scope of these assumptions.

Safety has a direct welfare consequence for the choice between title and a separate licence.

Proposition 4 (Safe title with separate licences). *Suppose Assumptions 1 and 2 hold, μ is supported on F , and g^S is F -safe. Under both original rights and g^S , every selected incumbent has retention probability one. Title and a separate licence therefore induce the same first-best effort*

$$e_i^{FB} \in \arg \max_{e \in E_i} \{b_i(e) - k_i(e)\}.$$

The two instruments generate the same terminal allocation in every supported state. Relative to original rights with optimized licences, vesting S raises reduced welfare by

$$\sum_{i \in V(S)} [v_i(1) - [v_i(1) - d_i]_+] = \sum_{i \in V(S)} \min\{v_i(1), d_i\}, \quad (3)$$

with zero allocation-option cost. Hence $S^(F)$ is an optimal support-safe pure settlement. It is the unique optimum if $\sum_{i \in C} \min\{v_i(1), d_i\} > 0$ for every $C \in S^*(F)$. When every licence is free, all support-safe title sets tie.*

Proof. The state-by-state part of Proposition 3 gives retention one under original and vested rights, so either instrument induces the displayed effort. Safety also gives state-by-state equality of complete allocations and therefore leaves every other retention and licence decision unchanged. Under original rights, selected user i 's optimized contribution is $[v_i(1) - d_i]_+$; under title it is $v_i(1)$. Their difference is $\min\{v_i(1), d_i\}$, proving (3). Cyclewise factorization makes every safe collection a subset of $S^*(F)$. Nonnegativity gives optimality, and strict positivity cycle by cycle gives uniqueness. $\square$

This proposition describes the zero-exchange-cost region of the planner’s problem. The gain per vested user is the licence cost where a licence would have been bought and the whole investment surplus where that cost would have deterred licensing. Thus title either saves a licence cost that would have been paid or restores investment deterred by that cost; with free licences the two instruments tie. Section 5 values unsafe title by deriving its retention insurance from future TTC paths; Section 5.4 locates the boundary between unique and redundant protection.

4.6 Implementation by periodic settlement

Periodic settlement separates the allocation clock from the rights clock. Fix exogenous settlement dates that partition time into epochs. Reference rights remain fixed within an epoch, apart from the append-only addition of each entrant’s own object. The institution may run TTC and implement its allocation for current use after any arrival. At the settlement date, it runs TTC from the accumulated reference record; that allocation becomes the rights state of the next epoch.

For a fixed final extension, every provisional clearing schedule and entrant order leads to the same settlement problem. Each TTC run remains individually rational relative to its reference rights and Pareto efficient in the active market. Successive settlement allocations are protected because each one becomes the next epoch’s reference-right vector. Online Appendix G states these properties formally and gives the proof.

Under unrestricted one-agent entry, one entrant separates any within-epoch change in an existing right (Proposition 2). Under finite support, Proposition 3 permits every certified cycle to vest before the common settlement date, and Theorem 2 lets this early-settlement set expand as the support shrinks. Periodic settlement is therefore the conservative baseline; the safety results identify the title changes that can occur sooner.

5 Investment and the Title-Insurance Hypergraph

In the safe region, title and a licence both deliver certain retention. We now turn to unsafe title, where changing one exchange right can insure investments elsewhere. We obtain an exact global representation on a stochastic single-gateway domain: valuable adaptation opportunities sit at designated targets, nontarget investment is normalized to zero, and every provisional target can hold a separate licence. These restrictions turn the general retention probability from Section 4.5 into a network object.

5.1 Stochastic interaction rings

The current TTC allocation partitions incumbents into nontrivial canonical cycles $\{C_q : q \in Q\}$. Write i^+ for the cyclic successor of i . In cycle C_q , agent i owns h_i and ranks

$$h_{i^+} \succ_i h_i \succ_i H \setminus \{h_i, h_{i^+}\};$$

the relative order within the last set is unrestricted. Current TTC assigns $x_i = h_{i^+}$. The institution chooses a vested set $V \subseteq Q$; write $Q \setminus V$ for the provisional cycles. It may give the target of any provisional cycle a separate licence; let $L \subseteq Q \setminus V$ denote those licensed targets and d_q the corresponding cost.

Each cycle has one designated target t_q . A state $\omega \in \mathcal{S}$ occurs with probability p_ω . At most one entrant 0_q is attached to each target in a state. She owns z_q , which t_q inserts above all old objects; every incumbent ranks every other new object below all old objects. Entrant 0_q requests h_{s_q} in cycle $C_{\sigma(q)}$ and ranks

$$h_{s_q} \succ_{0_q} z_q \succ_{0_q} (H \cup \{z_r : r \neq q\}) \setminus \{h_{s_q}\},$$

with arbitrary relative order in the last set. These ranking restrictions complete the stochastic single-gateway domain. In state ω the interaction graph has edge $q \rightarrow \sigma(q)$, so its outdegree is at most one; write $\mathcal{R}_\omega$ for its directed rings. For s, t in the same incumbent cycle, let $[s, t]$ be the directed segment beginning at s and following cyclic successors up to but excluding t , with $[s, s) = \emptyset$.

For a directed ring R in that graph, let $\rho(r)$ be the predecessor of r on the ring and define

$$E_{\omega R} = \{r \in R : s_{\rho(r)} \neq t_r^+\}, \quad A_{\omega R} = R \setminus E_{\omega R}.$$

The first set contains destinations whose incoming old-object demand is nonaligned. Such a demand survives only as long as the destination remains provisional, whereas an aligned demand survives vesting because the demanded object then belongs to the target who wants the entrant's object. Hence

$$R \text{ is live under } V \iff V \cap E_{\omega R} = \emptyset. \quad (4)$$

On a live ring, the entrant demands concatenate into one TTC cycle. The targets receive their new objects, the entrants receive the old objects they request, and the complementary incumbent segments unwind. A blocked or open interaction path, by contrast, restores the current incumbent cycles and leaves its entrants with their entry objects. The representation

theorem below states the exact allocation and its investment consequence: the vested set alone determines which rings survive, and hence which investors keep the objects they have adapted.

Let $B_{\omega R}$ be the allocation-welfare gain of a live ring relative to the current incumbent allocation. Up to a policy-independent constant, the allocation component of expected welfare is

$$A_\mu(V) = \sum_{\omega \in \mathcal{S}} p_\omega \sum_{R \in \mathcal{R}_\omega} B_{\omega R} \mathbf{1}\{V \cap E_{\omega R} = \emptyset\}. \quad (5)$$

The gain may be of either sign in the representation and in the unique- protection result below.

5.2 Retention hazards

Write $v_q = v_{t_q}$. A titled target can lie on a live ring only at an aligned destination, while a separately licensed target remains provisional and may also be exposed at a nonaligned destination. We therefore collect every state–ring pair containing q :

$$\mathcal{H}_q = \{(\omega, R) : R \in \mathcal{R}_\omega, q \in R\}.$$

The target’s retention probability under exchange-title set V is

$$\rho_q(V) = 1 - \sum_{(\omega, R) \in \mathcal{H}_q} p_\omega \mathbf{1}\{V \cap E_{\omega R} = \emptyset\}. \quad (6)$$

Within a state, directed rings are vertex-disjoint, so a target appears in at most one summand and the loss events are disjoint. Adding title can only block rings, so every ρ_q is nondecreasing. A hazard with an empty exposed set is uninsurable: no title blocks the ring that reassigns the adapted object.

Definition 7 (Authority-independent title-insurance hypergraph). For every (ω, R) and every $q \in R$, create a directed hyperedge

$$E_{\omega R} \rightrightarrows q$$

of weight p_ω . Its head is a potential investment exposed to the ring; its tail contains the exchange titles that can block that ring. A tail may contain its own head, and it may be empty.

A licence activates an investment head but leaves every tail unchanged; only exchange




Figure 1: The title-insurance hypergraph. Left: an interaction ring aligned at the investment head T_ℓ and nonaligned at the protector cycles P_u and P_v ; each edge points from the cycle receiving a new object to the cycle supplying the demanded right. A licence activates the head's investment and leaves the ring live. Title at either protector blocks the ring; title at the aligned head blocks nothing. Right: the induced hyperedge $\{P_u, P_v\} \rightrightarrows T_\ell$ of weight p_ω . Singleton tails and two-element tails form the tractability–hardness boundary in Theorem 4.

Alt text: Two-panel schematic. In the left panel, three cycle nodes form a directed interaction ring: the edge into the investment head is aligned; the edges into two protector cycles are nonaligned. In the right panel, the licensed investment head is linked by a two-element hyperedge to the two protector titles. Either protector title blocks the ring; title at the aligned head does not.

title can hit a tail and insure retention. Figure 1 displays the ring geometry that generates a two-protector hyperedge and the hyperedge it induces.

Theorem 3 (Interaction rings and the title-insurance hypergraph). *(i) For every state and vested set, entrant trade occurs exactly on the directed rings satisfying (4). On a live ring, target t_q receives z_q , entrant 0_q receives h_{s_q} , every incumbent in $[s_{\rho(q)}, t_q]$ keeps her current assignment, and the remaining nontargets in C_q receive their original objects. Off live rings, every incumbent keeps her current TTC assignment and each entrant keeps her entry object.*

(ii) The investment contribution of exchange-title set V and licence set $L \subseteq Q \setminus V$ is

$$J_\mu(V, L) = \sum_{q \in V \cup L} v_q(\rho_q(V)) - \sum_{q \in L} d_q. \quad (7)$$

(iii) Maximizing over licences gives the reduced planner objective

$$\tilde{\Psi}_d(V) = A_\mu(V) + \sum_{q \in Q} [\mathbf{1}\{q \in V\} v_q(\rho_q(V)) + \mathbf{1}\{q \notin V\} [v_q(\rho_q(V)) - d_q]_+]. \quad (8)$$

If every licence is free, all targets can be authorized and investment welfare is $\sum_q v_q(\rho_q(V))$. The title-insurance hypergraph remains a sufficient statistic.

Proof. Concatenate the paths created by a ring's entrant demands. At an unvested destination, $h_{s_{\rho(q)}}$ points to its original owner and incumbent pointers continue forward to target t_q . At a vested destination, liveness requires alignment, so the demanded object points directly to the target. The concatenation is a TTC cycle and gives the assignments in part (i). Removing it makes the complementary incumbent segments unwind to original objects. Every interaction path outside a live directed ring ends at a missing or blocked edge, or feeds a ring already removed. Backward induction then makes its entrants keep their new objects and restores the current incumbent cycles.

Part (i) assigns the entrant object to target q exactly when q lies on a live ring, so (6) gives retention. Title and licences determine whether v_q enters the objective; only V enters the live-ring indicators. Subtracting licence costs proves (7). For fixed V , a licence changes head q 's contribution from zero to $v_q(\rho_q(V)) - d_q$; optimizing head by head gives (8). At $d_q = 0$, nonnegativity makes authorization weakly optimal for every head. $\square$

5.3 Title or licence

The representation yields a marginal rule for the choice between title and a licence. Fix $q \notin V^-$, put $V^+ = V^- \cup \{q\}$, and hold fixed a set of authorized heads K satisfying $V^+ \subseteq K \subseteq Q$. Under the hybrid policy, q has a separate licence; under title, its licence is unnecessary. Define

$$O_q(V^-) = A_\mu(V^-) - A_\mu(V^+), \quad (9)$$

$$G_q(K, V^-) = \sum_{r \in K} [v_r(\rho_r(V^+)) - v_r(\rho_r(V^-))]. \quad (10)$$

The first term is the signed allocation-option cost of title; it is nonnegative when every policy-dependent ring gain is nonnegative. The second is the retention insurance supplied to all authorized heads.

Proposition 5 (Title versus separate licences). *(i) Holding K fixed, the welfare difference between giving q title and licensing it under V^- is*

$$W(\text{title at } q) - W(\text{licence at } q) = d_q + G_q(K, V^-) - O_q(V^-). \quad (11)$$

Title is weakly preferred exactly when $G_q(K, V^-) + d_q \geq O_q(V^-)$.

(ii) After all licences are reoptimized, define for $r \neq q$

$$\phi_r^{V^-}(\rho) = \begin{cases} v_r(\rho), & r \in V^-, \\ [v_r(\rho) - d_r]_+, & r \notin V^-, \end{cases}$$

and

$$\begin{aligned} \widehat{G}_q(V^-) &:= v_q(\rho_q(V^+)) - [v_q(\rho_q(V^-)) - d_q]_+ \\ &\quad + \sum_{r \neq q} \left[\phi_r^{V^-}(\rho_r(V^+)) - \phi_r^{V^-}(\rho_r(V^-)) \right]. \end{aligned} \quad (12)$$

Then

$$\widetilde{\Psi}_d(V^+) - \widetilde{\Psi}_d(V^-) = \widehat{G}_q(V^-) - O_q(V^-). \quad (13)$$

Hence title is weakly preferred in the reduced problem exactly when $\widehat{G}_q(V^-) \geq O_q(V^-)$.

If every $B_{\omega R} \geq 0$, then $O_q(V^-)$, $G_q(K, V^-)$, and $\widehat{G}_q(V^-)$ are nonnegative.

Proof. For part (i), let $D = \sum_{r \in K \setminus V^+} d_r$ be the licence cost common to both policies. Title produces

$$A_\mu(V^+) + \sum_{r \in K} v_r(\rho_r(V^+)) - D,$$

and the hybrid policy produces

$$A_\mu(V^-) + \sum_{r \in K} v_r(\rho_r(V^-)) - D - d_q.$$

Subtraction gives (11).

For part (ii), substitute V^+ and V^- into (8). The allocation difference is $-O_q(V^-)$. Head q 's optimized contribution difference is the first line of (12); every other head keeps the same title status, so its difference is the corresponding summand in the second line. This proves (13). Adding title only removes live rings, so every retention probability weakly rises. Monotonicity of v_r and $[v_r - d_r]_+$ makes all investment gains nonnegative. If allocation gains are nonnegative, blocking rings also gives $O_q \geq 0$. $\square$

A two-cycle margin. Let a costlessly licensed target q face one ring of probability λ , aligned at q and nonaligned at protector cycle p . If the ring is live it creates allocation gain B ; title at p also has outside-option cost γ . The licence permits investment and leaves retention

at $1 - \lambda$. Title at p is optimal exactly when

$$v_q(1) - v_q(1 - \lambda) \geq \lambda B + \gamma.$$

Although the protective title belongs to another cycle and need not authorize any valuable investment of its own, only vesting it restores full retention: its role is insurance, which the head's licence cannot reproduce.

5.4 Unique and redundant protection

The preceding analysis priced the choice between title and a licence at a single cycle. The economically relevant boundary lies one level up: whether an authorized investment has one way to insure its return or several substitute protectors.

The title-insurance hypergraph has *unique protection* when every nonempty tail is a singleton. For target q , let

$$\lambda_{qe} = \sum_{(\omega, R) \in \mathcal{H}_q: E_{\omega R} = \{e\}} p_{\omega}, \quad \beta_q = 1 - \sum_{(\omega, R) \in \mathcal{H}_q} p_{\omega}.$$

Then retention is the modular index

$$\rho_q(V) = \beta_q + \sum_{e \in V} \lambda_{qe}. \quad (14)$$

Theorem 4 (Unique versus redundant protection). *(i) Suppose the hypergraph has unique protection; allocation gains may have either sign.*

- (a) *For every collection of licence costs $d_q \geq 0$, the reduced objective $\tilde{\Psi}_d$ is supermodular. Its maximizers form a lattice, and an optimum is computable in polynomial time by submodular-function minimization when the component values have polynomial-time value oracles.*
- (b) *If investment is quadratic and either every licence is free or licences can be omitted without loss because $d_q \geq v_q(1)$ for every q , one polynomial-size maximum-closure network, and hence one minimum s - t cut, computes an optimum.*

(ii) Tails of size two suffice for both substitution and hardness.

- (a) *If a costlessly authorized head q faces one hazard of probability $\lambda > 0$ with tail $\{a, b\}$, and $v_q(1) > v_q(1 - \lambda)$, then, writing $J_q^0(S) = v_q(\rho_q(S))$,*

$$J_q^0(\{a, b\}) - J_q^0(\{a\}) - J_q^0(\{b\}) + J_q^0(\emptyset) = v_q(1 - \lambda) - v_q(1) < 0.$$

(b) *Maximizing integrated allocation and investment welfare is NP-hard even when every investment head has a costless separate licence, every incumbent cycle has two agents, every state satisfies the single-gateway restriction, every policy-dependent allocation gain is nonnegative, direct waiting costs are zero, investment is quadratic and rank preserving, and every hyperedge has at most two protectors. The result continues to hold when terminal allocation welfare is strictly above no-entry current welfare in every state under every title policy.*

Proof. For part (i), write $x_e = \mathbf{1}\{e \in V\}$. For head q , separate the weight on its own title:

$$r = \beta_q + \sum_{e \neq q} \lambda_{qe} x_e, \quad \ell = \lambda_{qq}.$$

Its optimized contribution is

$$F_q(x) = x_q v_q(r + \ell) + (1 - x_q)[v_q(r) - d_q]_+.$$

Conditional on $x_q = 1$, this is a convex nondecreasing function of a nonnegative modular index. Conditional on $x_q = 0$, the scalar function $[v_q(r) - d_q]_+$ is also convex and nondecreasing. Hence every cross-difference between two title coordinates distinct from q is nonnegative.

The marginal contribution of setting $x_q = 1$ is

$$\begin{aligned} m_q(r) &= v_q(r + \ell) - [v_q(r) - d_q]_+ \\ &= [v_q(r + \ell) - v_q(r)] + \min\{v_q(r), d_q\}. \end{aligned}$$

Convexity makes the bracketed increment nondecreasing in r , and monotonicity of v_q does the same for the second term. The cross-difference between q and every other title is therefore nonnegative. Summing over heads proves supermodularity of optimized investment welfare.

Under unique protection, each $E_{\omega R}$ is empty or a singleton. Its contribution to A_μ is therefore either the constant $p_\omega B_{\omega R}$ or the modular term $p_\omega B_{\omega R}(1 - x_e)$ for some title e . Thus A_μ is modular for allocation gains of either sign. Adding it to optimized investment welfare proves that $\tilde{\Psi}_d$ is supermodular and that its maximizers form a lattice. Minimizing its negative proves part (i)(a) (Iwata et al., 2001).

With free licences, $F_q = v_q(\rho_q(V))$. With no licences, $F_q = x_q v_q(\rho_q(V))$. Under quadratic effort, either expression expands into polynomially many nonnegative title monomials. Represent each monomial by a benefit node requiring its title variables. Each singleton- tail allocation term is a constant plus a signed linear title weight, which can be absorbed into the corresponding title node; empty-tail terms are constant. The resulting maximum-closure

network is polynomial in the explicit hypergraph size, and one minimum cut proves part (i)(b).

For part (ii)(a), either protector gives retention one; with neither, retention is $1 - \lambda$. Costless authorization keeps the investment active under all four title policies. Substitution gives the displayed negative cross-difference.

For part (ii)(b), reduce minimum vertex cover ([Garey and Johnson, 1979](#)). Create one protector cycle P_v for each vertex and one investment cycle T_ℓ for each edge $\ell = \{u, v\}$. In the edge state, a single-gateway interaction ring has aligned head T_ℓ and nonaligned destinations P_u, P_v , producing

$$\{P_u, P_v\} \Rightarrow T_\ell.$$

Separate mutually exclusive states make vesting each protector forgo one unit of allocation value. Scale quadratic investment so that covering an edge creates two units. After constants, welfare is

$$2|\{\ell : S \cap \ell \neq \emptyset\}| - |S|.$$

An uncovered edge can always be covered for a strict gain, so every optimum is a vertex cover; among covers the objective selects one of minimum size. Give every T_ℓ target a free licence. Because the incoming demand at T_ℓ is aligned, its title changes no state allocation or retention, so every policy has a welfare-equivalent policy that leaves all T_ℓ cycles provisional. Canonical TTC utilities implement every term, and a policy-independent aligned trade makes the state-by-policy allocation gain strictly positive. Online Appendix D gives the complete reduction. $\square$

Part (i) holds uniformly in the licence cost. Licence adoption may change endogenously as title raises retention, and the licence can be arbitrarily cheap. Even so, unique protection preserves increasing differences: titles are complements, and vesting one strengthens the case for vesting another. For interior licence costs the theorem delivers polynomial submodular minimization; the sharper one-cut representation applies to the free-licence and no-licence cases.

Part (ii) shows what breaks. The hard instances are otherwise benign: nonnegative trade surplus, single gateways, canonical two-agent cycles, investment values separable across heads, and protection supplied by titles the investor need not hold. Redundancy alone produces hardness: each protector insures several licensed investments and either endpoint can insure a given investment, so titles become substitutes and the planner faces a covering problem over exchange rights.

6 Information and Contingent Protection

The hypergraph also identifies when public information should change the timing of exchange title. We isolate one costlessly licensed investment head, which makes the comparison exact; several heads couple the protector choices and can restore the rank-two hardness.

6.1 One head and four information regimes

Fix a provisionally held investment head q with a costless separate licence. Every interaction ring below is aligned at q , so replacing the licence by head title would leave its allocation and retention unchanged. A public mode $j \in \mathcal{M}$ occurs with probability $\pi_j > 0$ and identifies one possible insurance edge

$$E_j \rightrightarrows q, \quad \emptyset \neq E_j \subseteq \mathcal{P}.$$

Conditional on mode j , its interaction ring arrives with probability λ_j . If it arrives and no title in E_j has vested, the target loses her current object and the ring creates allocation gain $B_j \geq 0$. If a protector vests, the ring is blocked.

Vesting protector e has additive conditional option cost $\gamma_e \geq 0$. The cost is incurred whenever e vests, including when the institution chooses title after observing the public mode. It is the reduced-form value of other allocation opportunities that title blocks; the mode signal leaves those opportunities unrevealed and intact. Let

$$\gamma(S) = \sum_{e \in S} \gamma_e, \quad \tau_j = \min_{e \in E_j} \gamma_e.$$

The mode signal arrives before protector settlement and effort, though it conveys no information about whether the ring itself will arrive.

We compare no disclosure; effort-only disclosure, where the title set is fixed before the mode while effort can respond; protector disclosure, where both title and effort respond; and hazard removal, which prevents every interaction ring but also destroys its allocation gain. Denote their optimal welfare by W^0 , W^{eff} , W^{disc} , and W^{rem} . Define

$$u_j = v(1 - \lambda_j) + \lambda_j B_j, \quad g_j = v(1) - u_j.$$

The net hazard loss g_j is the investment loss from the live ring minus its allocation gain.

Proposition 6 (Information, finality, and hazard removal). *In the one-head environment,*

$$W^0 = \max_{S \subseteq \mathcal{P}} \left\{ v \left(1 - \sum_{j: S \cap E_j = \emptyset} \pi_j \lambda_j \right) + \sum_{j: S \cap E_j = \emptyset} \pi_j \lambda_j B_j - \gamma(S) \right\}, \quad (15)$$

$$W^{\text{eff}} = \max_{S \subseteq \mathcal{P}} \left\{ \sum_{j: S \cap E_j \neq \emptyset} \pi_j v(1) + \sum_{j: S \cap E_j = \emptyset} \pi_j u_j - \gamma(S) \right\}, \quad (16)$$

$$W^{\text{disc}} = \sum_j \pi_j \max\{v(1) - \tau_j, u_j\}, \quad W^{\text{rem}} = v(1). \quad (17)$$

Moreover,

$$W^0 \leq W^{\text{eff}} \leq W^{\text{disc}}, \quad W^{\text{rem}} - W^{\text{disc}} = \sum_j \pi_j \min\{\tau_j, g_j\}. \quad (18)$$

After mode j is disclosed, protection is optimal when $\tau_j \leq g_j$, no protection is optimal when $\tau_j \geq g_j$, and both are optimal at equality.

Proof. For a fixed ex ante protector set, retention fails only when an unhit mode is realized and its ring arrives, giving (15). Conditional effort choice gives (16). Convexity and Jensen's inequality imply that effort-only disclosure weakly improves each fixed-set policy. When settlement also responds, choices separate by mode: protection yields $v(1) - \tau_j$ and no protection yields u_j . A contingent institution can always replicate a common set, proving the second inequality. Finally,

$$\max\{v(1) - \tau_j, u_j\} = v(1) - \min\{\tau_j, g_j\}.$$

Averaging gives the identity and the modewise protection rule. $\square$

The first gap in (18) is the productive value of information: effort adapts to posterior retention. The second is the legal value: the institution adapts which right becomes final. Although the weak value-of-information inequalities are standard, the decomposition and the hypergraph policy rule are the economic content.

The removal comparison turns on the signs of the net hazard losses. If every $g_j \geq 0$, costless prevention weakly dominates disclosure because it attains retention without a title cost. If every $g_j \leq 0$, disclosure instead weakly dominates blanket prevention because it leaves allocation-beneficial rings unblocked. With mixed signs, the expectation in (18) determines the ranking.

6.2 Fragmented protection

Let there be $m \geq 2$ equiprobable modes with distinct singleton tails $E_j = \{e_j\}$. Every title costs $c > 0$, every ring arrives with probability λ , and every live ring has allocation gain B . Write

$$g = v(1) - v(1 - \lambda) - \lambda B.$$

Corollary 1 (Information-induced finality). *In the symmetric fragmented environment,*

$$W^0 = W^{\text{eff}} = v(1) - \min\{g, mc\}, \quad W^{\text{disc}} = v(1) - \min\{g, c\}, \quad W^{\text{rem}} = v(1).$$

The value of protector disclosure is

$$\mathcal{I}_m(g) = \begin{cases} 0, & g \leq c, \\ g - c, & c < g < mc, \\ (m-1)c, & g \geq mc. \end{cases} \quad (19)$$

Without disclosure, no interior number of titles is ever a strict optimum: the institution vests none when $g < mc$ and all m when $g > mc$. After disclosure, it vests none when $g < c$ and only the realized title when $g > c$.

Proof. If k titles vest before the mode is known, welfare is

$$F(k) = v\left(1 - \lambda + \frac{\lambda k}{m}\right) + \lambda B\left(1 - \frac{k}{m}\right) - kc.$$

Convexity of v makes $F(k)$ convex, so a maximum occurs at $k = 0$ or $k = m$. The endpoint values are $v(1) - g$ and $v(1) - mc$. Under effort-only disclosure, welfare is linear in k with the same endpoint values, so $W^{\text{eff}} = W^0$. After protector disclosure, the institution compares $v(1) - c$ with $v(1) - g$ in each mode. Subtraction gives (19). $\square$

The region $c < g < mc$ is the central case: one relevant title is worth its cost, yet blanket ex ante protection is wasteful. The signal changes the legal policy from none to one even though its productive value is exactly zero. Information targets finality.

Symmetry is essential for the all-or-none statement. If $v(\rho) = \rho^2/2$, $\lambda = 1$, $B = 0$, mode probabilities are $(0.9, 0.1)$, and each title costs 0.1, no protection yields zero, protecting only the high-probability mode yields 0.305, and protecting both yields 0.3. Partial protection is uniquely optimal; however, the general formulas in Proposition 6 remain valid.

Disclosure versus prevention. If $g > 0$, Corollary 1 implies that disclosure captures the following fraction of the gain that hazard removal delivers over no disclosure:

$$\frac{W^{\text{disc}} - W^0}{W^{\text{rem}} - W^0} = \begin{cases} 0, & 0 < g \leq c, \\ 1 - c/g, & c < g < mc, \\ 1 - 1/m, & g \geq mc. \end{cases}$$

Indeed, Corollary 1 gives $W^{\text{rem}} - W^0 = \min\{g, mc\}$ and $W^{\text{disc}} - W^0 = \mathcal{I}_m(g)$. Dividing the three cases in (19) by the first expression gives the displayed fractions. If waiting for the mode signal costs $h \geq 0$, waiting replaces W^0 by $W^{\text{disc}} - h$, so it is weakly optimal if and only if $h \leq \mathcal{I}_m(g)$.

At high stakes, ignorance requires m titles, disclosure requires one, and prevention requires none. The value of waiting for information saturates at the option cost of the $m - 1$ titles that disclosure makes unnecessary.

6.3 The hypergraph recourse gap

Set $B_j = 0$. Let

$$\tau(\mathcal{E}) = \min\{\gamma(S) : S \cap E_j \neq \emptyset \text{ for every } j\}, \quad \bar{\tau} = \sum_j \pi_j \tau_j.$$

Call the environment a strong full-insurance region if a minimum-cost hitting set maximizes the effort-only objective and $v(1) - v(1 - \lambda_j) \geq \tau_j$ for every mode.

Corollary 2 (The cost of noncontingent finality). *In a strong full-insurance region,*

$$W^0 = W^{\text{eff}} = v(1) - \tau(\mathcal{E}), \quad W^{\text{disc}} = v(1) - \bar{\tau}, \quad W^{\text{rem}} = v(1).$$

Hence legal information is worth

$$\tau(\mathcal{E}) - \bar{\tau},$$

and, when $\tau(\mathcal{E}) > 0$, captures the fraction $1 - \bar{\tau}/\tau(\mathcal{E})$ of the hazard-removal benchmark.

Proof. By assumption, a cheapest hitting set maximizes the effort-only objective and yields $W^{\text{eff}} = v(1) - \tau(\mathcal{E})$. For every fixed title set, Jensen's inequality gives no-disclosure welfare weakly below its effort-only counterpart. A hitting set gives certain retention, so it attains the same value in the two regimes. Hence it also maximizes the no-disclosure objective and $W^0 = W^{\text{eff}} = v(1) - \tau(\mathcal{E})$. In disclosed mode j , the full-insurance inequality makes

protection weakly optimal, and its cheapest implementation is a tail member of cost τ_j . Thus $W^{\text{disc}} = v(1) - \bar{\tau}$. For any ex ante hitting set S , $\tau_j \leq \gamma(S)$ in every mode; averaging and then minimizing over S gives $\bar{\tau} \leq \tau(\mathcal{E})$. The formulas and differences follow. $\square$

The statistic is a recourse value: the premium for committing to titles before the institution learns which hazard it faces. Tail overlap makes a common protector useful across modes and reduces the gap, whereas fragmented authority enlarges it. For m disjoint singleton tails of cost c , it equals $(m - 1)c$.

The hyperedges, retention losses, and allocation gains in the information environment have an exact TTC implementation. Use one two-agent investment cycle and one two-agent protector cycle for each tail member; put the head on an aligned destination of the interaction ring and each protector on a nonaligned destination. Choosing cardinal utility gaps above the maximal adaptation return preserves all TTC rankings. The γ_e 's enter as primitives, standing for allocation opportunities outside the modeled interaction rings. Online Appendix E details the construction, states this convention, and records the computational checks.

7 Institutional Interpretation and Domains

The clause-level record comes from Welwyn Hatfield Borough Council (WHBC), the English social-housing authority whose decant policy appeared in the introduction. Its published policies document the timeline that the model formalizes. During a temporary decant, a tenant may use b and retain tenancy rights and payment obligations at permanent home a (Welwyn Hatfield Borough Council, 2025a, secs. 13.1 and 13.13). The tenant signs a licence for the temporary property, not a tenancy of it (Welwyn Hatfield Borough Council, 2025a, sec. 19.1). If the council recommends a permanent decant, the recommendation goes to the Exceptional Circumstances Panel. Approval leads to creation of a new tenancy at b , and the new tenancy is signed before the permanent decant is finalized (Welwyn Hatfield Borough Council, 2025a, secs. 19.2 and 19.4). In the paper's notation,

$$\underbrace{\text{temporary licence and use at } b}_{x_i=b, g_i=a} \longrightarrow \underbrace{\text{Panel approval}}_{\text{authorization to create title}} \longrightarrow \underbrace{\text{new tenancy at } b \text{ signed}}_{\tilde{g}_i=b}.$$

We take the signed new tenancy as the observable counterpart of vesting, with occupation and Panel approval as the use and authorization stages that precede it. The temporary licence records a distinct legal status, and the model treats investment authority as a separate instrument.

WHBC's secure-tenancy terms illustrate why title can carry several powers without making

them technologically inseparable. They provide a qualified exchange right and, for periodic secure tenants, a qualified right to make improvements, alterations, and additions subject to permission and other approvals ([Welwyn Hatfield Borough Council, 2025c](#), clauses 8.4 and 8.7). Sections 97–99 of the Housing Act 1985 govern consent to improvements for covered secure tenancies and state that consent is not to be unreasonably withheld ([UK Parliament, 1985](#)). A mutual exchange transfers tenancies together with their homes and requires a deed; surrender and re-grant is the alternative route ([Welwyn Hatfield Borough Council, 2025b](#), sec. 4.1). The exchangeable object is the tenancy-home record, while adaptation permission may be administered separately.

Assumption 1 treats the content and cost of the decant licence as primitives through d_i , and the main results cover $d_i = 0$. WHBC thus motivates the distinction among use, permission, and title, and the title-insurance results hold even when permission is free.

WHBC records the legal statuses and their timing. Assignment and exchange are administered case by case: the initial decant is landlord-directed, and mutual exchange is subject to eligibility, inspection, consent, and statutory refusal grounds. TTC supplies a transparent exchange technology in which the trading-option effect of rights can be characterized exactly. The portable policy content is the separation itself: record current use and reference rights separately, and identify the legal event that moves the second record.

Other housing institutions separate these margins at a broader scale. Households in Peru and Buenos Aires occupied their plots before formal title ([Field, 2007](#); [Galiani and Schar-grotsky, 2010](#)). China’s work-unit housing reform granted occupation, use, and inheritance under partial property rights while delaying resale and retaining a work-unit priority claim ([Wang, 2011, 2012](#)). WHBC provides the clause-level assignment and rights timeline used to map the distinction between current possession and transferable rights into the model.

The two theoretical blocks operate on different domains. The safety results allow arbitrary strict incumbent rankings and every entrant report on a finite object support. The title-insurance representation specializes to canonical current cycles, one designated target with a valuable investment opportunity per cycle, and stochastic single-gateway interaction graphs. Excluding nontarget investment and rerouting among competing targets yields an exact global retention formula for titled and licensed heads. The rank-two hardness shows that the resulting domain still contains the full covering problem.

Investment is relationship-specific, has no salvage value to its maker after reassignment, and preserves TTC rankings. Entry, preference rankings, the information process, and the legal instruments are exogenous; the planner adds cardinal allocation and investment welfare without transfers; and licence costs are separable across users, with no economies of scope in permissions. Both instruments authorize the same action and leave bargaining shares fixed.

Their welfare difference comes from the retention distribution generated by future exchange.

Proposition 6 and Corollary 1 hold for a single investment head, and the all-or-none policy before disclosure additionally uses symmetric fragmented protection; the asymmetric example in Section 6 locates that second boundary. With several simultaneous heads, protector choices couple across modes, and the covering structure of Theorem 4 governs them.

Safe title spends no exchange option. Relative to optimized licensing, it saves the licence cost when a licence would be purchased and restores the investment surplus when that cost would deter licensing. It ties a free licence. Unsafe title, however, can be insurance: its value depends on which future paths it blocks and which licensed investments those paths expose. Information is valuable when it arrives early enough to target that insurance.

8 Conclusion

An assignment, an investment licence, and exchange title perform different jobs. The assignment determines current use. The licence permits adaptation without changing the next market's endowment. Title changes that endowment and can therefore alter the probability that an investment remains with its maker.

Protecting assignments as they are implemented ratchets away future trades. The finite-support analysis identifies title that spends no exchange option. Safety factorizes across current TTC cycles and expands as uncertainty resolves. On a safe cycle, title and a separate licence induce the same allocation, retention, and first-best effort. Relative to optimized licensing, title saves the licence cost when a licence would be purchased and restores the investment surplus when that cost would deter licensing. It ties a free licence.

Outside the safe region, title is chosen when its retention insurance and saved licence cost cover the trade it blocks. The title-insurance hypergraph encodes this comparison. Unique insurance tails preserve supermodularity for every licence cost; two protectors, by contrast, are substitutes and make the problem NP-hard even when every investment head can be licensed for free. Information has a legal value because it identifies which of those titles is needed.

The timing lesson concerns legal records. Permission can precede exchange rights, and safe rights can move immediately. A trading option extinguished by an unnecessary title cannot be recovered. Unbundling authority enlarges the instrument set while leaving finality as a separate design problem.

Conflict of Interest

The author declares no conflict of interest.

Data and Code Availability

This article uses no primary or secondary data. Python code used for the computational verification exercises is included in the submission replication package, together with software requirements and reproduction commands. The code will be deposited in an approved public repository upon acceptance.

References

- Abdulkadiroğlu, A. and Sönmez, T. (1999). House allocation with existing tenants. *Journal of Economic Theory*, 88(2):233–260.
- Andersson, T., Ehlers, L., and Svensson, L.-G. (2016). Transferring ownership of public housing to existing tenants: A market design approach. *Journal of Economic Theory*, 165:643–671.
- Ashlagi, I., Liu, E., and Qian, P. (2026). Priority and externalities in dynamic matching. Working paper, <https://doi.org/10.2139/ssrn.7020858>.
- Ashlagi, I., Monachou, F., and Nikzad, A. (2025). Optimal allocation via waitlists: Simplicity through information design. *Review of Economic Studies*, 92(1):40–68.
- Balbuzanov, I. and Kotowski, M. H. (2019). Endowments, exclusion, and exchange. *Econometrica*, 87(5):1663–1692.
- Balbuzanov, I. and Kotowski, M. H. (2024). The property rights theory of production networks. *Theoretical Economics*, 19(4):1619–1658.
- Besley, T. (1995). Property rights and investment incentives: Theory and evidence from Ghana. *Journal of Political Economy*, 103(5):903–937.
- Blackwell, D. (1953). Equivalent comparisons of experiments. *Annals of Mathematical Statistics*, 24(2):265–272.
- Bloch, F. and Houy, N. (2012). Optimal assignment of durable objects to successive agents. *Economic Theory*, 51(1):13–33.
- Combe, J., Nora, V., and Tercieux, O. (2025). Dynamic assignment without money: Optimality of spot mechanisms. *Theoretical Economics*, 20(1):255–301.

- Doval, L. (2025). Dynamic matching. In *Handbook of the Economics of Matching*, volume 2, chapter 6, pages 489–577. Elsevier.
- Dworczak, P. and Muir, E. V. (2025). A mechanism-design approach to property rights. Conditionally accepted at *Econometrica*, <https://ellenmuir.net/papers/Dworczak-Muir-MDPR-2025.pdf>.
- Ehlers, L., Klaus, B., and Pápai, S. (2002). Strategy-proofness and population-monotonicity for house allocation problems. *Journal of Mathematical Economics*, 38(3):329–339.
- Ferguson, B. A. and Milgrom, P. (2026). Market design for surface water. In Lo, I. Y., Ostrovsky, M., and Pathak, P. A., editors, *New Directions in Market Design*, chapter 3, pages 65–87. University of Chicago Press.
- Field, E. (2007). Entitled to work: Urban property rights and labor supply in Peru. *Quarterly Journal of Economics*, 122(4):1561–1602.
- Galiani, S. and Schargrodsky, E. (2010). Property rights for the poor: Effects of land titling. *Journal of Public Economics*, 94(9–10):700–729.
- Garey, M. R. and Johnson, D. S. (1979). *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman, San Francisco.
- Grossman, S. J. and Hart, O. D. (1986). The costs and benefits of ownership: A theory of vertical and lateral integration. *Journal of Political Economy*, 94(4):691–719.
- Hafalir, I. E., Kojima, F., and Yenmez, M. B. (2025). Market design with distributional objectives: Efficiency, incentives, and property rights. *Journal of Economic Theory*, 230:106099.
- Hart, O. and Moore, J. (1990). Property rights and the nature of the firm. *Journal of Political Economy*, 98(6):1119–1158.
- Huitfeldt, I., Marone, V., and Waldinger, D. (2026). Designing dynamic reassignment mechanisms: Evidence from GP allocation. *American Economic Review*, 116(8):3036–3075.
- Iwata, S., Fleischer, L., and Fujishige, S. (2001). A combinatorial strongly polynomial algorithm for minimizing submodular functions. *Journal of the ACM*, 48(4):761–777.
- Kumar, R., Manocha, K., and Ortega, J. (2022). On the integration of Shapley–Scarf markets. *Journal of Mathematical Economics*, 100:102637.
- Kurino, M. (2014). House allocation with overlapping generations. *American Economic Journal: Microeconomics*, 6(1):258–289.
- Ma, J. (1994). Strategy-proofness and the strict core in a market with indivisibilities. *International Journal of Game Theory*, 23(1):75–83.
- Matsui, A. and Murakami, M. (2022). Deferred acceptance algorithm with retrade. *Mathematical Social Sciences*, 120:50–65.

- Nicolò, A., Salmaso, P., and Saulle, R. (2025). Dynamic one-sided matching. Working paper, <https://doi.org/10.2139/ssrn.5171244>.
- Nöldeke, G. and Samuelson, L. (2024). Investment and competitive matching. In *Handbook of the Economics of Matching*, volume 1, pages 125–222. Elsevier.
- 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.
- Shapley, L. S. and Scarf, H. (1974). On cores and indivisibility. *Journal of Mathematical Economics*, 1(1):23–37.
- UK Parliament (1985). Housing Act 1985. <https://www.legislation.gov.uk/ukpga/1985/68>. Sections 97–99.
- Wang, S.-Y. (2011). State misallocation and housing prices: Theory and evidence from China. *American Economic Review*, 101(5):2081–2107.
- Wang, S.-Y. (2012). Credit constraints, job mobility, and entrepreneurship: Evidence from a property reform in China. *Review of Economics and Statistics*, 94(2):532–551.
- Welwyn Hatfield Borough Council (2025a). Decant policy. <https://www.welhat.gov.uk/policies/decant-policy>. Effective March 2025.
- Welwyn Hatfield Borough Council (2025b). Mutual exchange policy. <https://www.welhat.gov.uk/mutual-exchange>. Effective September 2025.
- Welwyn Hatfield Borough Council (2025c). Tenancy agreement for secure tenancies. <https://www.welhat.gov.uk/downloads/file/1593/tenancy-agreement-for-secure-tenancies-booklet>.
- Weyl, E. G. and Zhang, A. L. (2022). Depreciating licenses. *American Economic Journal: Economic Policy*, 14(3):422–448.