

Recorded contracting after university license termination

Bolero

Abstract

Ending one university invention license does not necessarily end the invention's recorded contracting history. This study asks what agreement and financial records follow an eligible early termination, and which histories can support the narrower description of a strict recorded gap. We register a finite-record, non-causal analysis of APTO disclosure–agreement links. For each eligible disclosure, the index agreement is the one with the earliest qualifying recorded termination, with deterministic identifier tie breaking. Time zero is that agreement's termination date. A later agreement's effective date is its event timestamp; agreements beginning on or before termination cannot witness post-termination contracting. The primary outcome is a commercial agreement with disjoint recorded party sets within three calendar years after termination. Conservative logical bounds retain incomplete follow-up in the denominator. Selected analyses examine continuity, one-year positive classified receipts at distinct-agreement grain, and a five-year extension. The eligible cohort is *[pending]*, with primary logical range *[pending]*. The registered interpretation is *[pending]*. This is a disclosed amendment informed by an operator's review of a prior-run manuscript. No empirical results were supplied to this author, and all empirical slots remain unfilled.

Keywords: University technology transfer; license termination; subsequent contracting; incomplete follow-up; logical bounds.

PRE-DATA DRAFT: empirical placeholders are not findings.

Introduction

University technology transfer is a sequence of agreements, development decisions and recorded payments. A terminated license is an informative event in that sequence, but its interpretation depends on what follows and on what the records can observe. A later contract may document renewed interest, continuation of an existing arrangement, or administrative succession. An absent later contract may reflect either a complete record without a qualifying event or insufficient observation. Distinguishing these

cases matters for interpreting a university’s portfolio without turning administrative records into claims about technological or social success.

The empirical question is: among APTO disclosures with an eligible recorded early license termination, how often does a subsequent commercial agreement with disjoint recorded parties begin within three years after termination; which witnessed histories contain a strict recorded gap; and do the distinct subsequent agreements show a positive classified receipt within their own first year? The object is a currently recorded disclosure history, not the causal effect of ending a license. We retain financial outcomes alongside contracting while separating their populations, clocks and observation requirements.

Prior work already studies licensing, termination, commercialization and royalties. Shane’s *Selling University Technology: Patterns from MIT* develops a transaction-cost account of these outcomes. The Dechenaux and coauthors literature examines development and termination under uncertainty. Those precedents make an expansive novelty claim inappropriate. Our narrower contribution is the explicit separation of eligible terminated-disclosure histories, later recorded counterparties, strict continuity guards and bounded financial observation. Incomplete full-text access prevents a claim that no previous study follows comparable sequences. [Shane \(2002\)](#); [Dechenaux et al., SIEPR version \(2006\)](#).

This manuscript accompanies an amended registration. An operator’s inspection of a prior-run manuscript motivated correction of the origin from agreement effectiveness to termination. The design author received no empirical results, but that separation does not make the amendment an untouched outcome-blind registration. The seven scientific units, their horizons and registered interpretation rules remain fixed. Amendment history and engineering checks are retained in the accompanying provenance records; no empirical finding is supplied by those checks.

Institutional reasoning and competing explanations

A licensee’s development decision can depend on uncertain technical prospects, financing, demand and the value of continued development. A subsequent contract therefore has several possible explanations. Another organization may undertake further work; a parallel license may remain active; or an existing commercial relationship may be documented under another agreement. The records supplied here cannot discriminate these explanations causally. Nor does absence of a subsequent direct agreement establish that knowledge was never transferred.

Current institutional descriptions make renewed contracting plausible without establishing its historical frequency. UW CoMotion discusses seeking another developer when development stops. Michigan’s licensing description allows that a sublicensee may obtain a direct university license after termination. These are contextual descriptions of institutional possibilities, not APTO coverage attestations, a universal legal rule, or policy treatment assignments. [UW licensing FAQs](#); [Michigan license process](#).

Financial consideration also differs from product sales. Upfront, maintenance and milestone fees can yield positive ledger receipts at different stages of development.

We therefore measure the existence of a positive classified inflow, without interpreting it as net revenue, profit, first sale or technology use. The binary outcome is deliberate: the APTO dictionary documents mixed revenue/cost events, incomplete date/classification support and a transaction ordinal that is not universally unique. A receipt total would require additional unsupported accounting decisions. The distinction also separates this design from the broader Stanford technology-transfer study identified in the literature search. [Liang et al. \(2022\)](#).

The underlying economic distinction is between terminating a particular contractual relationship and exhausting an invention’s possible uses. The inherited account in Jensen and Thursby’s *Proofs and Prototypes for Sale* connects early-stage licensing to inventor effort and incentives; the inherited Dechenaux, Thursby and Thursby contract-theory account distinguishes incentives from shelving. These arguments motivate why termination might be followed by a different arrangement, but do not determine which sequence will occur in this cohort. Their use here is limited to the rationale documented in the registration, not a newly verified empirical finding.

The closest termination research also gives a reason not to equate a missing subsequent agreement with failed knowledge transfer. The registered reading of the June 2006 Dechenaux and colleagues manuscript notes that a firm can end a license after learning knowledge that is not protectable. Conversely, a new direct agreement can formalize an existing relationship. Our incremental contribution is a longitudinal measurement framework that preserves these distinctions and makes unresolved histories visible, not a general ranking of universities or a causal assessment of licensing policies.

The broad measure establishes that a later-party record exists; continuity examines whether it represents a clean recorded gap; receipts add a different, agreement-level financial observation. The stricter chronology can qualify a renewal narrative, and payment information can qualify its financial content, without turning either into product success or welfare. The three questions remain interpretable when their evidence points in different directions.

Data and measurement

The authorized boundary contains fourteen APTO tables; this study uses six: agreement, disclosure, disclosure_agreement, agreement_counterparty, counterparty and agreement_revenue_event. The structural documentation lists ten university namespaces, but it establishes neither runtime availability nor empirical coverage. The governed planner must bind actual physical projections, versioned dictionaries, extraction identities and source metadata. No observed counts, missingness distributions, event ranges or identifying variation are assumed.

A disclosure can link to several agreements, and an agreement can cover several disclosures. IDs retain tenant prefixes. Entity joins must be many-to-one after identical projected duplicates collapse and conflicting entity keys are quarantined. Link tables become distinct pair sets. Counterparty sets are reduced before joining disclosure histories. Broken candidate links remain possible events, rather than disappearing through an inner join. Ledger events are not deduplicated on event_seq; the existential receipt indicator is invariant to repeated positive rows. Names, titles, abstracts, addresses and external registry links are not used.

Commercial-license classes are bound to a provider dictionary before outcome computation. A qualifying class explicitly grants commercial use or development rights to university intellectual property. Option-only, nondisclosure, material-transfer, service/research, administrative amendment-only and definitively noncommercial classes are excluded. Mixed or unresolved classes remain unknown. Unresolved separability excludes a namespace before the cohort is frozen, even if a tentative mapping lists a qualifying code; contradictory definitive metadata stops execution. No text, payment-size or outcome-based classification is allowed.

Cohort and the three distinct dates

For each disclosure d , eligible index candidates have known disclosure date no later than agreement effective date, effective date strictly before termination date, and termination strictly before known expiration. A missing expiration is allowed and flagged; a known signature after termination excludes that candidate. Invalid or unavailable index dates and unresolved index types do not qualify. The index is the first eligible *recorded termination*, not a first-ever termination or business failure.

Let $A(d)$ be that eligible agreement set, let $\tau(a)$ denote `termination_date(a)`, and let $\text{id}(a)$ denote the exact agreement identifier. The deterministic definition is

$$a_0(d) = \arg \min_{a \in A(d)} (\tau(a), \text{id}(a)),$$

$$t_0(d) = \text{termination_date}(a_0(d)).$$

The second coordinate means exact lexicographic identifier ordering, used only for ties. We separately retain `index_effective_date = effective_date(a0)`. For any later candidate a_1 , its event timestamp is $s(a_1) = \text{effective_date}(a_1)$. Index selection, termination origin and event timestamp are distinct operations. An agreement can begin after the index effective date yet before its termination; it is not a post-termination witness. All contracting anniversaries and the strict separation guard start at t_0 . Receipt follow-up instead starts at s .

Atom 1 assesses all ten namespaces using its fixed support table. Completed inspections can establish supported, evidenced-unsupported or unresolved index scope. A positive supported cohort may coexist with excluded unresolved namespaces; its denominator never represents their unknown populations. An incomplete required inspection is operational `no_result`. Once the supported cohort is frozen, N counts disclosures and remains unchanged for the primary, strict and five-year quantities. We also report distinct index-agreement counts. No complete-case or older-cohort restriction removes unresolved follow-up.

Broad later-party event and missing-information bounds

A definite witness requires a distinct qualifying agreement currently linked to the disclosure, complete nonempty recorded party sets with no common ID, and

$$t_0(d) < s(a_1) \leq H_h(d), \quad H_h(d) = \text{add_years}(t_0(d), h).$$

Calendar anniversaries replace February 29 by February 28 in a nonleap destination year. The event cannot have a known contradictory interval ending at or before

its effective date. Such a contradiction remains unresolved unless an independent necessary condition definitively rules out the candidate. Known common parties or a known out-of-window start rule out that path. Missing party sets, types, dates or valid links can leave it possible. A signature later than the horizon does not move a broad effective-date event; signature chronology is reported separately.

For each d and h , the lower indicator is one when a definite witness exists. The upper indicator is zero only when every candidate is decisively nonqualifying and metadata certifies the necessary histories and links through the anniversary, with a snapshot reaching it. Otherwise the upper indicator is one. A known witness is conclusive despite other unresolved candidates. At one, two and three years we apply the prespecified marginal monotonic closure: lower indicators accumulate forward and upper indicators take the minimum over later horizons. Contradictory closures stop execution.

$$[L_h, U_h] = \left[N^{-1} \sum_d \ell_{dh}, N^{-1} \sum_d u_{dh} \right].$$

These are conservative logical ranges for unresolved indicators in the finite cohort. They are not confidence intervals and are not claimed to be jointly sharp. Shared agreements can couple unresolved facts. The bounds do not address misclassified witnesses, excluded undated terminations, unrecorded disclosures, corporate lineage or historical membership of intellectual property in a contract. No Kaplan–Meier extrapolation, regression, standard error or p-value is registered.

Continuity, strict gaps and receipts

Atom 3 examines all broad witness paths within three years. A strict path additionally requires distinct known agreement-local IDs, no known party-local-ID collision with complete local-ID support, a signature strictly after index termination and no later than the subsequent effective date, at least 31 civil days between termination and that effective date, no simultaneous eligible index termination, and supported absence of another commercial agreement overlapping the open gap. Active intervals are start-inclusive and end-exclusive, using the earlier known termination/expiration endpoint. An unknown end is not evidence of dormancy. The 31-day guard is an author-set bookkeeping-succession screen, not a legal threshold.

A definite failed requirement dominates missing requirements. A disclosure is strict if any broad path passes, non-strict if all broad paths definitely fail, and otherwise unresolved. If the diagnostic is non-strict, mixed or unresolved, Atom 4 applies the same stricter predicate to all possible paths in the original N , including disclosures without a definite broad witness. It reports strict bounds nested within the broad possibilities and a conservative range for broad events that fail strictness. Passing this screen never proves rights reversion, corporate independence or a newly recruited developer.

Atoms 5–6 use M distinct agreement IDs that are broad witnesses for any cohort disclosure within three years. Bundled disclosures do not multiply ledger outcomes. For agreement $a1$ the receipt window is inclusive from $s(a1)$ through $\text{add_years}(s(a1), 1)$. A positive receipt requires a dated finite amount greater than zero with provider

is_revenue equal to true. A later refund does not cancel this existential outcome. Fiscal year cannot substitute for event_date, and total_revenue is an unused hybrid cumulative scalar. A certified empty ledger can support zero; an uncertified empty ledger is unresolved. Atom 5 describes whether indicators are decidable. Atom 6 retains all M, including unresolved agreements, and reconverges financial support with any scientific completion of the continuity diagnostic.

Adaptive interpretation and longer follow-up

The primary three-year attention threshold is 10%. Equality belongs to the material lower-bound region. An upper bound strictly below 10% is low; otherwise the range straddles the threshold. Witness-positive and witness-absent states are distinguished because continuity and receipt analyses require a witnessed population. Low or unresolved primary states activate Atom 7, which extends the same snapshot and fixed cohort to five years after termination. It neither changes the three-year result nor follows newly found late agreements into receipts.

For each disclosure, the registered paired increment range is

$$[\Delta_d^L, \Delta_d^U] = [\max(0, \ell_{d5} - u_{d3}), \min(1, u_{d5} - \ell_{d3})].$$

Averaging these endpoints gives the outer range for a first broad effective-date event after year three and by year five. No optional candidate-level tightening is allowed. A late signature is not a late event. The extension threshold is five percentage points; it is an attention rule, not an economic-equivalence margin. All outcome-selected analyses are exploratory descriptive, with no holdout and no claim of independent confirmation. There are no confirmatory statistical tests and hence no statistical multiplicity family. Every branch retains its registered disposition, including not_activated, boundary, otherwise and no_result.

Registered dispatch policies and publication closeout

Standalone execution respects registered activation. The runtime may instead declare direct dispatch and a budget B before execution, selecting the first B Molecule members without waiting for outer activation. A direct attempt still needs genuine scientific inputs and the original conditional scope. Shared cohort, broad-event, continuity and receipt-support calculations can be computed within that attempt when predecessor records are absent. They are recorded as local evidence, not accepted predecessor World results. Existing failed predecessors are not silently retried. All scientific definitions and selection restrictions remain fixed. For example, all-strict continuity does not license the selected strict refinement, and a material primary state does not license the selected late extension; attempted out-of-scope branches receive otherwise with no numerical result.

The manuscript identifies locally constructed continuity separately from an accepted Atom 3 result. Such evidence cannot fill missing predecessor World conditions in composite claims. Final closeout records the dispatch policy, explicit activation bypass and the disposition of all seven members. Members outside a declared direct budget are budget_unselected; incomplete selected work prevents final closeout. Disclosure limits precede all presentation decisions. A closed consuming-Atom disposition

takes precedence over missing-source preview wording, including the receipt continuity slot. Pending is reserved for an incomplete preview. This V004 repair follows the supplied review’s four synthetic failing routes and the operator’s dispatch request; it changes no scientific estimand and does not claim new empirical validation.

Results

Cohort scope and recorded later contracting

The support and exclusion assessment is *[pending]*. The frozen cohort contains *[pending]* and *[pending]*. Figure 1 displays first exclusion reasons, with unidentifiable/conflicting key groups in a separate unit panel; Table 1 also reports composition by tenant, termination calendar year, source field and disclosure kind, with explicit missing categories. Composition is *[pending]*. These summaries describe the selected records and do not establish complete termination histories in APTO.




Fig. 1 The identity-defect panel counts quarantined key groups, not known inventions. The disclosure panel counts eligible disclosures and first exclusion reasons in frozen supported-index partitions. Index selection uses earliest eligible termination, with agreement-ID tie breaking. Time zero is that selected agreement’s `termination_date`. Cohort size: *[pending]*. Composition by tenant, termination year, field and disclosure kind appears in Table 1.

Table 1. Cohort construction and composition

Table 1 component	Count/status	Denominator or unit
Included namespace: source-qualified label	—	Namespace; label determined by approved source
Excluded namespace: source-qualified label and reason	—	Namespace; no missing population inferred
Eligible disclosures	—	Frozen N

Table 1 component	Count/status	Denominator or unit
Distinct selected index agreements	—	Unique agreement IDs
Tenant composition: approved category	—	N disclosures
Termination year: observed calendar category	—	N disclosures
Termination source field: source category	—	N disclosures
Disclosure kind: source category	—	N disclosures
Missing category within each composition dimension	—	N disclosures
Quarantined ambiguous/conflicting key groups	—	Key groups, separate from disclosure counts

Notes: the implementation expands the registered category rows rather than combining unrelated dimensions. Quarantined key groups are separate from known disclosures. Excluded scope is not represented in the reported logical ranges. Figure 1 carries the corresponding flow display.

At three years after index termination, the broad later-party interval is *[pending]*. Its witnessed, certified-negative and unresolved disclosure counts are *[pending]*. The one- and two-year ranges and the three-year uncertainty share are *[pending]*. Figure 2 uses later effective dates as event timestamps. Among definite witnesses, earliest known effective-date delays are *[pending]*. Signature chronology is *[pending]*; late or missing signatures do not shift the event into another anniversary bin.

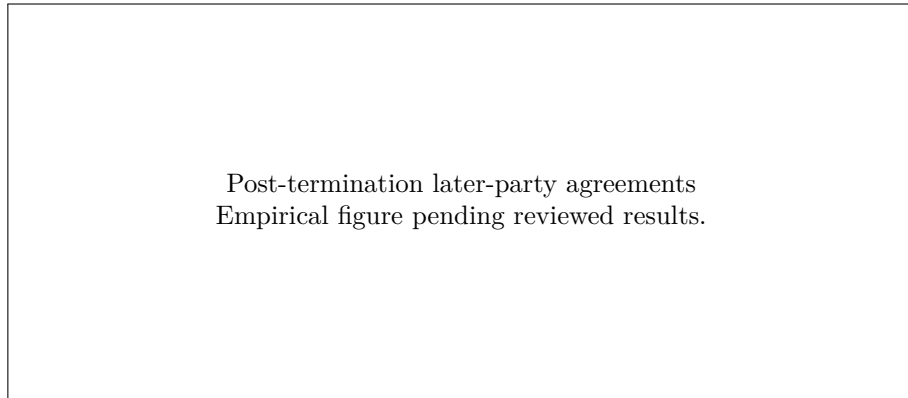



Fig. 2 Cumulative logical bounds at one, two and three anniversaries after index termination; they are not confidence intervals. Later agreement effective_date is the event timestamp. Starts on or before termination do not qualify. The timing panel conditions on a definite witness and displays its earliest known qualifying effective date. The dotted line is the author-set 10% attention threshold. Primary range: *[pending]*.

The primary interpretation is *[pending]*. If the lower bound reaches 10%, at least one in ten eligible recorded disclosures has a qualifying later-party record within three years. If the upper bound is below 10%, the full registered logical range is below that author-set threshold, with no equivalence inference. If the interval straddles 10%, it does not classify the share relative to the threshold. Zero witnesses with a positive upper bound leave subsequent contracting unresolved. A source boundary or operational failure leaves this quantity unavailable; it cannot be replaced by a numerical zero or a selected follow-up result.

Continuity and the stricter outcome

The continuity diagnostic is *[pending]*, and the path assessment is *[pending]* (Figure 3; Table 2). First-reason buckets combine failed and unresolved requirements under neutral predicate labels; the strict-path bucket means all requirements pass. The nonexclusive flag table separately preserves passes, failures and unknowns. Thus an unknown earlier end can support possible overlap without establishing an overlapping agreement. Its registered interpretation is *[pending]*. If every witnessed disclosure has a strict path, that statement concerns the witnessed population alone; it does not establish equality of full-cohort broad and strict bounds. A mixture, no strict paths, or unresolved strictness prevents an unqualified gap/re-entry narrative.



Strict-gap diagnostic among witnessed disclosures
Empirical figure pending reviewed results.

Fig. 3 Witnessed disclosures are classified by whether any path passes every strict recorded-gap requirement, all paths fail, or strictness is unknown. Path-level reasons use their own denominator. The open gap begins after index termination; it never begins at index effective_date. Diagnostic: *[pending]*. First-reason buckets combine failed and unresolved requirements under neutral predicate labels; the strict-path bucket means all requirements pass. The nonexclusive flag table separately reports passes, failures and unknowns, so possible overlap is never stated as observed overlap.

The conditional full-cohort strict analysis is *[pending]*, with broad-minus-strict range *[pending]* (Figure 4). If activated, its bounds use the same N and termination origin as the primary analysis. If it is not activated, its slot reports that disposition; no strict estimate is inferred from nonactivation. Any source boundary leaves the broad analysis intact within its original scope.




Fig. 4 Broad and strict intervals use the same frozen disclosure denominator and the third anniversary of index termination. The strict rule adds identity, signature, 31-day separation, overlap and retained-history guards. No legal rights reversion follows from passing. Selected interval: *[pending]*. The dotted line marks the author-set 10% attention threshold.

Table 2. Later contracting and continuity

		Certified nega- Witnessed		Unresolved	Upper	
Table 2A: broad horizon	N	W	Z	Q	Lower L	U
1 year	—	—	—	—	—	—
2 years	—	—	—	—	—	—
3 years	—	—	—	—	—	—

Table 2B: continuity and stricter scope	Count/denominator	Lower bound	Upper bound
Witnessed disclosure population W3	—	Not a range	Not a range
Strict S among W3	—	Not a range	Not a range
Non-strict F among W3	—	Not a range	Not a range
Unresolved V among W3	—	Not a range	Not a range
Strict share among witnessed W3	W3	—	—
Selected strict full-cohort quantity	N	—	—
Broad-minus-strict outer range	N	—	—
Full-cohort strict witnesses / unresolved	—	Not a range	Not a range

Table 2C: existing reason/flag evidence	Assessment	Numerator	Denominator/unit
First-reason category from registered output	—	—	—
Nonexclusive predicate flag from registered output	—	—	—
Signature chronology category from registered output	—	—	—
Known effective-date delay interval	Observed interval	—	Witnessed disclosure count

Bindings: `annual_bounds` and `broad_counts` supply 2A; `continuity_counts`, `strict_interval` and `strict_difference` supply 2B; `continuity_reasons`, `signature_chronology` and `contract_delays` supply 2C. Source categories are data-dependent labels, not an invitation to create a grouping. All original column names are in `results-map.md`. Notes: intervals are logical uncertainty ranges, not confidence intervals. First-reason and nonexclusive counts cannot be summed together. The conditional strict rows retain nonactivation or boundary wording when appropriate; the witnessed-population diagnostic never supplies a missing full-cohort estimate.

Financial records and the longer horizon

Receipt observability among distinct witnessed agreements is *[pending]* (Figure 5). Conditional financial outcomes are *[pending]*, with earliest known dated receipt timing *[pending]* (Figure 6; Table 3). The financial interpretation is *[pending]*. Positive witnesses establish classified inflows. A certified zero concerns only the registered one-year ledger window. No definite receipt with a positive upper bound is unresolved evidence. A financial measurement boundary or failed read is not evidence of zero financial value. If completed inspection establishes that the frozen ledger contract changed and its indicators cannot be reproduced, the receipt result is a source boundary; an unexplained mismatch is operational `no_result`. Neither route replaces the frozen agreement population or yields a refreshed financial rate. The reported continuity context is *[pending]* and never selects a favorable receipt subsample. This copied context is withheld when its source continuity state is disclosure-limited, even if financial quantities are independently releasable.

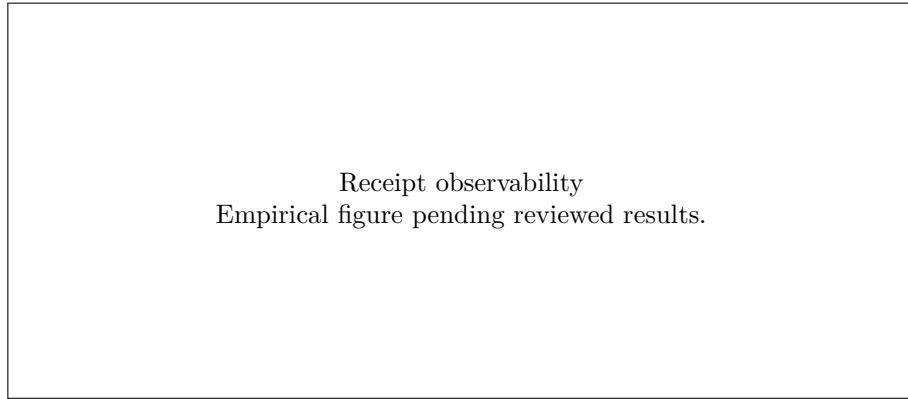


Fig. 5 Support for the one-year receipt indicator among distinct later agreement IDs witnessed within three years after index termination. Each receipt window instead begins at that later agreement's effective_date. The support split does not report positive versus zero receipts. Support: *[pending]*.

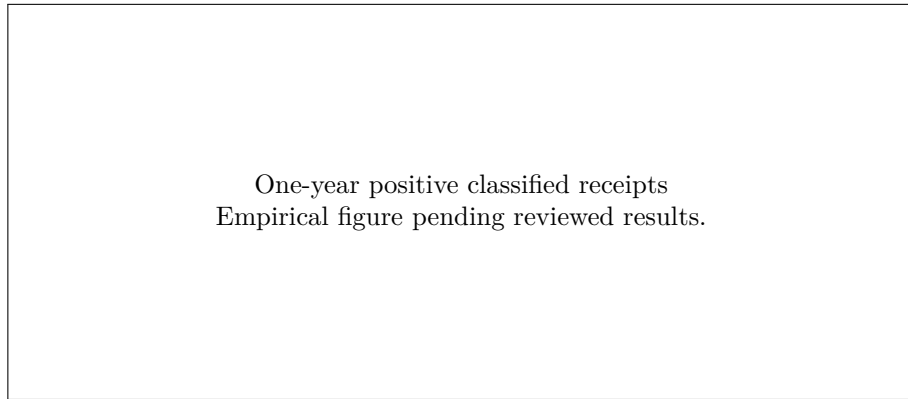


Fig. 6 Logical bounds for ever-positive classified ledger inflow within the inclusive year beginning on the later agreement's effective_date. This receipt origin differs from the index termination origin for contracting. Timing is conditional on a known dated positive receipt; undated potential receipts can precede it. No netting, revenue sum or sales inference is made. Receipt range: *[pending]*.

Table 3. Receipt observability and realization

Table 3 component	Count	Denominator	Lower	Upper
Distinct witnessed agreements M	—	Agreement universe	Not a range	Not a range
Decidable receipt indicators K	—	M	Not a range	Not a range
Unresolved receipt indicators	—	M	Not a range	Not a range

Table 3 component	Count	Denominator	Lower	Upper
Positive dated receipts	—	M	Not a range	Not a range
Certified zero receipts	—	M	Not a range	Not a range
Positive-receipt logical range	Not a count	M	—	—
Earliest known	—	Known-	Not a range	Not a range
positive-receipt delay, registered interval		positive agreements		
Continuity context	Source- qualified state	No receipt subgroup	Not a range	Not a range

Notes: positive means a finite dated amount greater than zero classified as revenue. Repeated positive ledger rows do not multiply this existential outcome. Certified zero, unresolved ledger history and unavailable computation are different dispositions. The table does not report monetary totals, net returns or product sales.

The selected five-year cumulative range is *[pending]* and the paired increment is *[pending]* (Figure 7; Table 4). Its interpretation is *[pending]*. A lower increment of at least five percentage points establishes delayed first effective-date events under the registered outer bound. An upper bound below that threshold describes a limited increment through year five; a straddling range remains unresolved. Nonactivation supplies no five-year finding. The primary three-year conclusion remains the primary conclusion in every branch.




Fig. 7 Selected same-cohort cumulative ranges at years three and five after index termination and the paired outer bound for a first event in between. The event timestamp remains the later agreement’s effective_date; a late signature never creates a late broad event. Shared unresolved facts can make these outer bounds conservative. The lower panel dotted line marks the author-set five percentage-point threshold. Increment: *[pending]*.

Table 4. Selected five-year extension

Table 4: fixed-cohort extension				
	Denominator	Lower	Upper	Unresolved
Three-year broad interval	N	—	—	—
Five-year broad interval	N	—	—	—
Paired first-event increment, years 3–5	N	—	—	—

Notes: this table is conditional under the registered policy. A direct attempt outside its scientific condition yields its registered nonnumerical disposition. Late signatures do not create late effective-date events, and unavailable follow-up is never set to zero.

Discussion

The combined registered interpretation is *[pending]*. The central distinction is between an observed contractual continuation and a claim about economic recovery. A witnessed later-party agreement documents a sequence in retrospective records. A strict recorded gap provides a narrower chronological description. A positive classified receipt adds financial evidence at agreement grain. These quantities answer different questions and need not yield a single favorable narrative.

When later agreements are witnessed but strict continuity is absent or uncertain, the appropriate account is subsequent contracting with a restricted gap interpretation. When strict paths are established, they still do not identify corporate independence, legal scope or treatment effects. When receipts are observed, they establish inflow rather than profitability or product development. When observation is insufficient, the corresponding question remains open. Conditional findings cannot repair excluded index scope or replace the primary interval.

The study cannot identify a causal termination effect. Termination, later contracting and financial observation are endogenous to invention prospects, partner capability, demand and recording. No instrument, discontinuity or policy change is documented. Scheduled expiration is not random assignment, and additional controls would not supply the missing identification. The finite-record design makes the observable question explicit without converting temporal order into causality.

Several limitations persist even with flawless execution. Current links need not describe contract membership at signature; distinct party IDs need not indicate independent corporations; unknown dates and unretained histories can widen ranges; and selection into disclosure, eligible termination and supported namespaces limits generalization. Revenue sums, profit, products, patents’ downstream impact, employment and welfare are not supported outcomes here. The six authorized projections supply no product-sales, development-effort, ownership-history or net-cost accounting measure. These omissions are schema boundaries, not negative findings. Financial scope is retained through the supported binary receipt question.

Reporting, access and conclusion

Confidential row-level derivations and source locators remain in governed evidence. Candidate tables and figures require the existing disclosure assessment, including complementary suppression; a ten-unit minimum is the registration’s conservative floor, and stricter provider rules prevail. Small-cell state labels can themselves disclose information and may require a nonnumerical public disposition. The release disposition is *[pending]*. Final closeout accounts for all seven registered members; a partial preview is explicitly labeled and cannot stand in for a complete execution account. No actual transaction example, confidential identifier or prior-run empirical manuscript is reproduced here.

The supplied analysis and figure code can be checked before data access. These checks verify selected algorithms and physical geometry on artificial inputs; they do not establish data availability, population completeness or scientific acceptance. Funding, affiliations and provider publication-review rights have not been supplied and remain explicitly unresolved disclosure items. The package compiler retains the manuscript class and owns final formatting and PDF validation.

The study’s eventual conclusion must retain the distinction between post-termination contracting, strict recorded gaps and receipts. Its evidentiary reach is determined by the registered records, observation bounds and composable interpretation rules. The final scoped conclusion is *[pending]*.

Supplementary methods

Relational construction and denominator preservation

The input sequence first resolves entity keys, then forms distinct disclosure–agreement and agreement–party links, and only then constructs histories. Identical projected entity rows collapse; conflicting entity identities are quarantined. Reducing party sets before joining histories avoids multiplying agreement rows by the number of parties. The subsequent disclosure-level indicators are existential: several qualifying agreements still count as one witnessed disclosure. In the receipt analysis those same histories yield a distinct agreement roster, so a bundled agreement does not contribute once for every covered disclosure.

Incomplete links have scientific consequences. A broken candidate link or unknown party set can leave a later event possible even when no definite witness is recovered. Dropping such rows before applying the predicate would turn a measurement problem into an apparently negative history. The lower endpoint therefore uses definite witnesses; the upper endpoint excludes an event only when necessary conditions definitively fail and the required observation evidence is present. A fully observed negative history and a recent or incompletely retained history remain different cases.

Index support and follow-up support also operate at different stages. A namespace lacking a usable index definition is outside the frozen cohort, with its unresolved scope reported. Once the cohort is frozen, incomplete later observation stays inside its denominator as uncertainty. This is why the full-cohort estimates cannot be reinterpreted as complete-case rates or rates for every invention in APTO. Table 1 reports the first distinction; Tables 2–4 preserve the second.

Event chronology and strictness

An index agreement’s effective date can substantially precede its termination. The index is selected by the earliest eligible termination, whereas a later candidate is timed by its own effective date. The three operations are not interchangeable. An agreement effective before or on the index termination date is not a post-termination witness, regardless of when its signature appears. Anniversary arithmetic uses calendar years, with the registered February-29 convention, rather than a fixed number of elapsed days.

Strictness is assessed on every broad witness path. It adds local-ID checks, signature ordering, the 31-day separation screen, simultaneous-index checks and the absence of a conflicting active agreement during the open gap. Active intervals include their start and exclude their end. A known failed requirement is decisive even if another requirement is unknown; only an entirely passing path establishes strictness. Conversely, a missing earlier endpoint does not establish a dormant gap. First-reason buckets provide a compact display, while the nonexclusive flags retain exactly which conditions pass, fail or remain unknown.

The gap restriction is a bookkeeping and chronology screen. Its 31-day value is not a legal boundary and does not establish return of rights, independent corporate ownership or recruitment of a new developer. These are distinct constructs for which the necessary evidence is not supplied.

Bounds, attention rules and sensitivity

At each horizon the registered indicators yield exact fractions with the frozen N. The monotonic closure accumulates definite lower witnesses forward and uses later upper bounds to constrain earlier horizons. An inconsistent closure is not smoothed away. Dependence through shared agreements can make the displayed ranges conservative; no sharpness or independent-Bernoulli precision claim follows.

The primary 10% and extension five-percentage-point cutoffs are attention rules. Their registered sensitivity quantities describe how many currently unresolved indicators would have to resolve in a particular direction to change the classification. The calculation retains both the required unresolved count and the available unresolved count, including an explicit unreachable case. These quantities diagnose dependence of an interpretation on missing information; they are not p-values, posterior probabilities or a new equivalence test. The primary result and exact fraction are retained regardless of classification.

Financial observation and composite interpretation

A positive receipt can be established by a qualifying observed ledger event. A negative receipt indicator requires the stronger observation statement that the relevant classified history is complete for the agreement’s one-year window. Neither an empty delivered ledger nor an undated potential payment independently supplies that statement. A refund does not erase the existence of an earlier positive inflow under this definition; its existence also does not establish positive net revenue.

The paper combines these analyses through their registered interpretation, not by counting positive Atom labels. Broad contracting can remain interpretable when strictness is unresolved. A receipt boundary leaves separately supported contracting evidence intact. The selected long-horizon analysis cannot replace the primary three-year statement. Mixed evidence therefore narrows particular claims rather than forcing the whole study into a favorable or unfavorable category. All seven scientific units retain a reported disposition in *[pending]*.

References and source scope

Shane, S. (2002). Selling University Technology: Patterns from MIT. *Management Science*, 48(1), 122–137. DOI: [10.1287/mnsc.48.1.122.14281](https://doi.org/10.1287/mnsc.48.1.122.14281).

Jensen, R., and Thursby, M. (2001). Proofs and Prototypes for Sale: The Licensing of University Inventions. *American Economic Review*, 91(1), 240–259. DOI: [10.1257/aer.91.1.240](https://doi.org/10.1257/aer.91.1.240). The theoretical use is inherited from the registration’s abstract-level source account.

Dechenaux, E., Goldfarb, B., Shane, S., and Thursby, M. Appropriability and Commercialization: Evidence from MIT Inventions. The registered detailed passages refer to the June 2006 SIEPR manuscript; the final article is identified by DOI [10.1287/mnsc.1070.0780](https://doi.org/10.1287/mnsc.1070.0780). Version-specific observations are not attributed to an unread final-journal passage.

Dechenaux, Thursby and Thursby (2009), contract-theory source retained in the identification document. DOI: [10.1016/j.ijindorg.2008.05.001](https://doi.org/10.1016/j.ijindorg.2008.05.001). Its inherited incentive/shelving rationale is used without a new full-text-access claim.

Liang, Elrod, McFarland and Zou (2022). Systematic analysis of 50 years of Stanford University technology transfer and commercialization. *Patterns*, 3(9), 100584. DOI: [10.1016/j.patter.2022.100584](https://doi.org/10.1016/j.patter.2022.100584). Institutional sources are linked in the text. Exact retrieval scope and limitations are retained in the accompanying literature record.