

Fund Design as a Computational Problem

AnchorPoint Labs

October 2, 2026

ABSTRACT

Private fund structuring requires simultaneous consideration of legal constraints, investor characteristics, investment activities and commercial objectives. Experienced counsel assess these interactions, but the cost of expert time limits how many alternatives a sponsor can examine before selecting a design. This paper describes a computational approach to exploring the design space of a fund: the structures that satisfy the governing constraints for a stated set of facts, and the consequences that distinguish them. The approach represents a fund as connected entities and participants, applies characterized legal rules to confirmed facts, and maps the candidate structures within the model's coverage and their consequences on a common basis. An illustrative buyout fund shows what exploring that space reveals: two responses that each address a new investor class's exposure but differ in who bears the cost, how much participation flexibility they preserve and how they hold up as the fund evolves; a structure whose diagram is unchanged while its consequences change materially; and a design that is preferable under expected conditions but fragile under plausible ones. The paper presents the framework, states the division of labor between characterized rules and language models that follows from it, applies it to the illustrative fund and closes with a description of Structura, a fund architecture platform built on these principles.

Keywords: private fund formation; fund structuring; computational design; constraint modeling; design space exploration; structural alpha; language models

Introduction

A semiconductor is simulated before fabrication. A building is modeled before construction. In both cases, computational tools let professionals test alternatives against interacting constraints before committing resources; in chip design, checking a layout against explicit design rules before fabrication dates to the 1970s [1]. Engineers call the set of designs that satisfy the governing constraints the design space, and its systematic examination design space exploration; the view of design as a search through a space of alternatives under constraints goes back to Simon [2]. Private fund structuring presents a comparable problem, although its constraints arise from legal rules and commercial requirements rather than physical laws, and they depend on interpretation, evolving authority and facts that may still be uncertain when design begins. The approach described here is a first-principles one: the design is derived from the governing constraints themselves, with

legal rules taking the place of physical laws, rather than from what was built last time.

A fund is designed before it is drafted. The manager describes the strategy and likely investors, and fund counsel proposes an architecture: a limited partnership, a general partner, a management company and, where appropriate, feeders or parallel vehicles. Tax and regulatory specialists assess the consequences for the fund, its manager and each investor class. Of the work involved in forming a fund, this structuring decision is likely the most consequential. The documents implement an architecture that has already been chosen, and a drafting error can often be corrected by amendment. A structure that is wrong for its investors, that lacks the flexibility the fund later needs or that carries avoidable tax cost is usually lived with, because reorganizing entities after capital has been committed carries tax, consent and cost consequences of its own.

The economics of that work are beginning to shift. Language models are lowering the cost of producing documents, and drafting has been the practical barrier to customization: a structure that departed from precedent required agreements that departed with it. As that barrier falls, demand for structures tailored to a particular investor base, strategy or jurisdiction is likely to grow, and each must be designed before it can be drafted. Tooling for the drafting step is arriving quickly. Comparable tooling for the design step, where the alternatives are mapped and a proposed structure is compared with them, has been largely absent. That is the problem this paper addresses.

Computational design is the approach this paper proposes for that step. Starting from stated facts and explicit constraints, a model can generate the candidate structures it knows how to represent, identify the consequences that distinguish them and expose the assumptions behind them, so that the alternatives are in view before a structure is selected. The modeled design space, used in that sense throughout, is the portion of the design space within the model's coverage. Section 1 describes the current state of fund structuring and why most of the design space goes unexplored. Sections 2 through 4 describe the framework, including where language models belong in it and where they do not. Section 5 applies it to an illustrative fund. Section 7 collects the limits of the approach in one place, and Section 9 states which of the described functions Structura implements today.

1 The current state of fund structuring

The constraint on fund structuring as practiced today is the cost of exploring the design space. Each alternative architecture requires further expert time. Should certain investors participate through a blocker or a parallel vehicle? Does a sovereign investor change the analysis? What happens if the strategy adds loan origination? A manager may examine only a few. The alternatives left unexplored are not necessarily worse; their merits are simply unknown.

That cost has two sources. The first is scarcity of skill. Structuring is the most skill-intensive part of fund formation, because it requires holding the tax, regulatory, ERISA and commercial consequences of an architecture in mind at once and knowing from experience which of them will matter in practice. It is learned late in a lawyer's training and in most firms is partner-level work, priced accordingly. Court filings show partner rates at the

largest firms running from roughly \$1,200 to nearly \$2,500 an hour and senior associate rates approaching \$1,500 as early as 2023 [3], and reporting in 2026 puts the top partner rates above \$3,000 with partner rates at the fifty largest firms up 16% in a single year [4]. As an illustration, a structuring call that brings together a corporate partner, a tax partner and a senior associate at those rates costs the sponsor on the order of \$5,000 to \$7,000 an hour before any further specialist joins. At half those rates the constraint is the same in kind: each additional alternative requires further expert time, and that time is the most expensive on the engagement. Designing from first principles for a novel strategy or investor base is therefore realistic mainly for the largest sponsors, and even there investors have grown wary of formation expenses they ultimately bear: the Institutional Limited Partners Association has pressed for limits on organizational expenses since 2019 [5], and its 2026 guidance describes those expenses as continuing to escalate, with legal fees the primary driver [6]. The second is fragmentation. A design passes through several lawyers and specialists, each responsible for part of the analysis. At the time of design, that division creates gaps: where each assumes that another has examined an interaction, an issue can go unexamined. Afterward, it leaves the analysis dispersed across the people who performed it, their correspondence and the documents that implemented it, so that reconstructing it, let alone re-running it when a fact changes, is expensive. The process is artisanal in the way that baking bread is artisanal: it produces excellent results in experienced hands, and it does not scale well when each order is different.

Precedent is the rational response to both. Most new funds begin from a structure that worked before, drawn from the standard architectures the practitioner treats as a catalogue [7]. Precedents preserve accumulated experience and provide a practical starting point, but a familiar architecture can conceal assumptions about investor composition, income character or the location of activities that no longer hold, and reliance on it is the reason most of the design space remains unexplored.

The sequence of the process compounds the problem. That sequence is not a failing of counsel; it is market practice, and it is a way of containing the initial cost, since fewer people in the room for fewer hours is the least expensive route to a workable structure. The first-order decisions, the vehicle type, the jurisdictions and the participation path for each investor class, and the second-order decisions that follow from them are therefore generally made within the law firm, shaped by its

precedents and by the preferences those precedents carry. The structure is presented to the client, which may include the marketing team and the placement agent, when it is already largely settled. Marketers know what the market wants: which investors are in the pipeline, what terms they will accept and what vehicles they have refused before. Lawyers know how to implement it. When the two meet only after the architecture is fixed, the design is built without the demand information, and the demand side receives it as a finished product, a diagram and a memorandum. Every further question is a legitimate request for more expert time, so questions are rationed: asking what a different investor mix would do, or why one vehicle was chosen over another, competes with the budget for the work itself. The same economics that made the sequence efficient are what keep the people closest to the market from interrogating the design, and reconciling the two sides afterward is expensive because the structure is by then embedded in documents.

2 The structure as an analytical object

2.1 Relationships determine consequences

A fund structure can be represented as a directed graph. Its nodes represent investors, legal entities and portfolio holdings. Its connections represent relationships such as investment participation, management and the movement of capital or compensation. Each element carries attributes relevant to the analysis, including jurisdiction, legal role and tax classification.

This representation matters because consequences depend on the relationships among elements. The presence of a corporation somewhere in a diagram does not establish that a particular investor's exposure has been addressed. The analysis must follow that investor's participation through the relevant entities to the underlying activity. The same entity can therefore have different significance depending on where it sits and whose participation passes through it.

A strategy label is similarly incomplete. Two funds described as private equity funds may hold different types of assets, use different forms of financing and admit different investors. Their legal and economic consequences can differ despite a shared commercial description. A useful model needs enough factual detail to distinguish those cases without treating every descriptive difference as material.

The graph provides a common object for legal, cost and economic analysis, so that a proposed change is examined

across all three on the same set of relationships rather than on assumptions that differ from one part of the analysis to another.

2.2 Characterizing legal rules

Under this approach, legal rules are characterized so that they can be applied to the represented structure and its stated facts. A finding identifies the relevant consequence, the part of the structure affected and the supporting authority. Findings derive from the facts presented by a design rather than from its name or resemblance to a standard diagram. The idea that legal rules can be expressed in a form that a machine can apply to facts has a substantial literature of its own, from the early formalization of statutes as logic programs [8] through computational law [9] to the "rules as code" work of public-sector rulemakers [10] and recent programming languages designed for legislation [11]; fund structuring is a bounded domain in which it can be put to practical use.

Characterization requires legal interpretation. Someone must determine which facts matter, when a rule applies and what conclusion follows. A rule may reflect settled authority, an interpretation or an issue on which the authority is still moving. The treatment of limited partners under the self-employment tax rules is a current example: the Second Circuit's affirmance of the Tax Court's functional analysis in September 2026 changed the confidence with which a common management-company arrangement could be recommended [12]. Those distinctions affect how a conclusion should be used and the review it requires, and an explicit representation makes an interpretation inspectable without making it correct. The supporting authority and the conditions under which it is applied must remain available for review, which is why each characterized rule should correspond to an identifiable provision or authority rather than to a paraphrase of several; the legal knowledge-systems literature calls this property isomorphism [13].

The role of language models in this work deserves a precise statement, because they are well suited to several parts of it and poorly suited to one. A language model is prediction software: it produces the text most likely to follow from its inputs. That makes it a capable assistant in building a rule library, where every proposal is reviewed before it is used. It can propose structured representations of statutes, regulations and cases, draft test cases that exercise those representations and flag inconsistencies among them; the lawyer determines the interpretation to encode, and the tests check that the implementation does

what the lawyer decided. It is similarly useful at intake, translating a manager’s description of a proposed fund, or an existing structure chart, into inputs that the user confirms, and asking for the facts the description leaves out rather than supplying them.

Designing the fund is a different kind of problem. A structural conclusion must follow from the confirmed facts and the applicable rules, must be the same each time those facts and rules are the same, and must be traceable to the authority that supports it. A prediction of the most plausible answer does not satisfy those requirements. Systematic studies have found that general-purpose language models produce incorrect legal statements at high rates when asked direct questions about the law [14], and that even purpose-built legal research tools that retrieve authority before answering still do so at a meaningful rate [15], and a fluent answer that cannot be reproduced or attributed is difficult for counsel to review. The approach described here therefore uses language models where prediction is the right tool and applies the characterized rules deterministically: the same confirmed inputs, rule library version and settings produce the same results. The dividing line is between proposing and concluding. A language model may propose a representation of a rule, a reading of a manager’s description or a change to a structure worth examining, and it may explain what the rules found. Every structural conclusion is produced by the characterized rules applied to the confirmed facts.

The point survives improvements in the models. Suppose a future model answered every legal question put to it correctly. It would still answer only the questions put to it, and in fund formation the difficulty is often knowing which questions a change raises. A venture manager that decides to add secondary purchases of portfolio-company shares to a fund’s mandate has changed nothing on the diagram: the same partnership, the same general partner, the same investors. To an untrained eye the structure is unaffected. The adviser, however, may have relied on the exemption from registration available to advisers solely to venture capital funds, which turns on the fund holding mostly qualifying investments acquired directly from portfolio companies, and secondary purchases are not qualifying investments beyond a limited basket [16]. Nobody would think to ask an oracle about adviser registration when the question on the table was investment strategy. A representation that carries the adviser’s status, the fund’s holdings and the rules linking them raises the question itself, because the change to the holdings is evaluated against every rule that reads them.

That scaffolding, the structure as an explicit object against which all the characterized rules are run, is what turns a correct answering machine into a design tool.

The same scaffolding is what makes a language model useful, and safe, in a third role: searching the design space. Given a structure, the facts behind it and the findings the rules produced, a model can propose changes that a practitioner would consider worth testing, such as interposing a corporation on one participation path, adding a feeder or a parallel vehicle, or removing a vehicle that appears to earn nothing. Each proposal is a hypothesis. The model does not evaluate it. The change is applied to a copy of the structure and the characterized rules, the cost model and the economic model are run on the result, so that the proposal is accepted or rejected on the same basis as every other candidate and takes its place in the same comparison. Three constraints keep that arrangement honest. The model may change the structure but never the facts: a proposal that would work if an investor were not tax-exempt is a question for the user, not an edit. The search is bounded, a fixed number of proposals in a fixed number of rounds, so that it reports what it tried and when it ran out of useful ideas rather than continuing indefinitely. And the model’s observations about a structure, that an arrangement is unusual, that a vehicle appears redundant or that an input looks inconsistent, are reported as observations, separately from findings, and cannot alter a finding, a candidate’s viability or a fact the user stated. A model used this way can criticize what the rules concluded. It cannot change it.

2.3 Facts and assumptions

A proposed fund is rarely fully specified when structuring begins. Commitments may be uncertain, prospective investors may have different requirements and the investment mandate may still be developing. The model must distinguish confirmed facts from assumptions so that an apparently precise result does not obscure an incomplete description. In practice that is a three-way distinction: facts the user has stated, facts the model has supplied by default from what was stated, and facts that are not known at all. A fact in the third category that could change a consequential conclusion is a question to put to the user, not a gap to fill by assumption, and that applies with particular force where a language model is helping with intake, since supplying a plausible fact is exactly what prediction software does when nothing stops it.

An assumption can be material even when changing it leaves the diagram intact. It may alter an investor's tax exposure, the adviser's obligations or the expected cost of operating the structure. Materiality therefore extends beyond whether an alternative architecture becomes preferable; it also concerns whether the consequences of the existing architecture change enough to require attention. Each unresolved fact is a dimension along which the design space can shift. The analytical task is to identify which of them deserve investigation before selection: facts capable of changing a consequential conclusion warrant more attention than details that leave the space unchanged.

3 Exploring the design space

The idea that legal planning can be treated as a search through a space of alternatives under legal constraints is not new. Schlobohm and McCarty described estate planning in those terms in 1989: a search over estate plans constructed from an explicit representation of the Code, starting from prototype plans and modifying them by deformations because searching every possible plan was impractical [17, 18]. McCarty later stated the general form, a search through a space of possible transaction patterns for the course of action that, under the rules of the legal analysis system, best achieves a given set of goals [19]. The lineage runs from TAXMAN in 1977 to a 2026 system that pairs a language model with a constraint solver, and it is set out in the note at the end of this paper.¹

This paper does not claim that idea. It applies it to the architecture of a private investment fund and extends it in the ways the rest of this section and Section 4 describe: the structure itself is represented as a graph, so that a vehicle has consequences because of where it sits on an investor's participation path; competing multi-entity architectures are generated and compared, rather than one proposed architecture being checked; several bodies of law are applied to the same structure; legal, cost and economic consequences are evaluated on a common basis and traced to the investor class that bears them; the comparison is re-run as the facts change; and the result is treated as a space to be explored rather than a plan to be delivered. The reference structures from which candidates are generated are prototypes in the sense of EPS II, and the bounded changes of Section 2.2 are deformations, proposed by a language model and tested by the rules.

3.1 Constraints and commercial preferences

The design space of a fund is bounded and shaped by requirements that operate differently. Some modeled conditions make a candidate inconsistent with its intended legal treatment; these define the boundary of the space, and a structure that fails one lies outside it. Others create an exposure that particular investors may accept, require the sponsor to address or regard as commercially unacceptable. Considerations such as administrative burden and flexibility require comparison rather than a pass or fail conclusion. Requirements of the second and third kinds do not draw the boundary; they shape the terrain within it.

These distinctions should remain visible. A commercially attractive feature cannot cure a failure to satisfy a necessary legal condition. Conversely, a finding that an investor faces an additional obligation does not by itself make the structure unavailable; its significance depends on the investor, the activity and the applicable requirements.

3.2 Multiple objectives

There may be no single point in the design space that is preferable on every relevant dimension; the problem is one of multiple objectives in the sense of the decision-analysis literature [20]. One arrangement may accommodate a wider investor base while introducing additional expense. Another may simplify operations while leaving less room for different participation terms. The design problem is to understand those exchanges well enough to choose deliberately.

A useful distinction is between an alternative that is inferior on every relevant dimension and one that offers a meaningful advantage at a cost. The former warrants little further attention. The latter is an alternative that cannot be improved on one dimension without loss on another, and it presents a genuine decision: where lower cost and greater flexibility point toward different structures, selecting between them requires a view about the fund's priorities.

That view may differ among participants. A sponsor may value the ability to admit future investors, while an initial investor may focus on the expenses charged to its capital. A structure that improves an aggregate outcome can still impose an unattractive burden on one class. Exploration therefore needs to identify who receives a benefit and who bears its cost.

3.3 Interacting decisions

Structural choices are interdependent. A change introduced to address one investor's exposure can create entity-level tax, additional expenses or operating requirements. A decision about where participation occurs may also affect allocation flexibility and the practical administration of investor rights. Evaluating each choice separately can miss its effect on the rest of the arrangement.

Interposing an entity, for example, may change the treatment of income reaching an investor. The analysis must account for the consequences within that entity as well as at the investor level, because an exposure may have been relocated or exchanged for a different cost rather than removed. The same reasoning applies to simplification: removing a vehicle reduces direct expenses while limiting the ability to treat particular investors or investments differently. Entity count is informative, but it is not a sufficient measure of design quality.

4 Economic consequences and durability

4.1 Consistent comparisons

Legal compatibility alone does not determine whether a structure is worth building. Formation expenses, recurring obligations and the allocation of economic outcomes all affect its suitability. The model associates cost and economic assumptions with the represented design so that alternatives can be examined on a common basis.

A meaningful comparison holds the underlying investment assumptions constant while examining what changes because of the structure. Otherwise an apparent structural advantage may reflect a different assumed portfolio, financing arrangement or investment result. The horizon and measurement basis must also be consistent: a formation expense cannot be compared directly with an annual cost or an undiscounted tax amount accumulated over the life of a fund.

The location of a modeled cost matters as much as its amount. Entity-level tax, investor-level tax and withholding have different economic effects, and a model must state which effects it includes. A reduction in one reported category should not be read as an equivalent improvement in an investor's total outcome without examining what occurs elsewhere. Simulating a distribution of investment outcomes can show how sensitive the comparison is to provider charges, the

timing and character of income, realizations and negotiated terms.

4.2 Sensitivity and structural change

Sensitivity analysis asks whether a conclusion survives a change in the facts [21]. Some changes make a structure gradually more or less attractive relative to the alternatives. Others redraw the boundary of the design space: they introduce a new constraint or make a previously unnecessary vehicle worth considering. Distinguishing those effects explains both the current choice and the conditions under which it should be revisited.

Fund size provides a simple example. Recurring expenses associated with an additional vehicle become easier to absorb as commitments increase, but the value of that vehicle also depends on which investors use it and what their participation requires; there is no general commitment threshold that establishes its suitability independently of those facts. Investor composition creates another kind of sensitivity. A new investor class may introduce an exposure that the existing design did not need to address, or the same architecture may continue to work while its expenses, obligations or modeled tax consequences change. Both outcomes matter to the sponsor, even though only the first suggests a different diagram.

4.3 Designing for an evolving fund

A structure is selected at a point in time but operates over an extended period. The investor base, portfolio and sponsor's business may develop in ways that were uncertain at launch. A useful exploration therefore considers whether a design remains suitable across plausible future conditions and what changing it would require.

Flexibility has value when it preserves a realistic option that would otherwise be costly or difficult to implement, a point the engineering-design literature makes in the language of real options [22]. It also has a carrying cost. Establishing vehicles for remote possibilities may impose years of expense without a corresponding benefit. The question is which future developments are sufficiently plausible and consequential to justify accommodating them now.

We use structural alpha to describe the potential value retained through better structural choices, that is, by selecting a better structure than the precedent would have supplied. That value may arise from avoiding

unnecessary expense, addressing tax exposure or preserving useful options. Establishing it requires a comparison with a credible alternative under consistent assumptions.

5 Exploring the design space of an illustrative fund

Consider a buyout fund investing in operating businesses held through entities treated as partnerships for U.S. tax purposes, with acquisition leverage at the fund level. Hold the portfolio, the leverage and the economic assumptions constant, and widen the investor base in stages. What follows is a qualitative account of the design space a model of this kind maps for that scenario and of how the space moves as the investor base widens. It reports the direction and location of the effects and the facts on which each conclusion turns, not the model's scores or the rules by which candidates are generated.

5.1 Domestic taxable investors

With only U.S. taxable investors, the design space is small and its best structure is easy to see. A single domestic limited partnership is the natural baseline, and the generated alternatives that add a feeder or an offshore vehicle fall behind it for a simple reason: they add entities and cost without any investor who needs them. The analytical contribution at this stage is to make that reasoning explicit. The single partnership is preferable because of the investor base, not because it is simple, and the exploration records the circumstances that would undermine it.

5.2 Adding tax-exempt investors: two responses to one exposure

Adding U.S. tax-exempt investors redraws the design space. The single partnership now carries two findings on the exempt investors' participation path: the operating income of the pass-through portfolio is unrelated business taxable income, and the fund-level acquisition debt brings the debt-financed income rules into play for income not already so characterized [23]. Two candidates emerge as the credible responses, and the interesting result is that they are close to each other on the aggregate measures while differing in ways that matter to particular participants.

The first is a single fund with a blocked feeder: the exempt investors subscribe to a feeder that invests through a U.S. corporation into the main fund. The second is a main fund with a parallel fund for the exempt

investors, the parallel fund investing through its own corporation alongside the main fund in each portfolio company. Both address the exposure the same way, by placing a corporation between the exempt investors and the pass-through portfolio, and both carry a similar amount of modeled tax leakage, because in each case the corporation pays entity-level tax on the income that passes through it.

In the two candidates as modeled, what separates them is not the tax result but three things around it. The first is cost and its allocation. The parallel structure adds an entity, and with it formation cost and recurring administration, audit and tax preparation. The blocked feeder is cheaper to form and cheaper to run. In both structures the corporate tax is borne by the exempt sleeve alone; the taxable investors' outcome is unchanged. An aggregate leakage figure for the fund therefore conceals that one class is paying all of it, which is exactly the kind of result an investor in that class would want to see stated.

The second is participation flexibility. The feeder as modeled invests in the main fund on the main fund's terms, so its investors share the main fund's investment path and allocation mechanics, and it blocks everything that passes through it; side letters can still give those investors differentiated rights, but not a different portfolio. The parallel fund as modeled is a separate partnership with its own agreement. It can carry different allocation provisions and excuse rights, and it can decide investment by investment whether to interpose its corporation, which matters when the portfolio will include companies held in corporate form and acquired without fund-level debt, which need no blocking at all. That flexibility is what the additional entity buys, and it is a feature of this parallel arrangement rather than of every parallel fund.

The third is suitability as the fund evolves. In the modeled scenario the choice turns on fund size: there is a commitment level below which the parallel fund's carrying cost outweighs the flexibility it preserves, and where that level falls depends on the cost assumptions and on how much weight the sponsor gives to flexibility. Removing the exempt capital altogether makes both candidates unnecessary and returns the comparison to the single partnership. Neither candidate is right in general. The parallel structure is favored by a larger fund, a meaningful and heterogeneous exempt sleeve, an expectation of corporate-form investments or of excuse and opt-out mechanics, and a sponsor who values differentiated terms. The blocked feeder is favored by a

smaller fund, a homogeneous exempt sleeve content to share the main fund's investment path, and a portfolio that will need blocking throughout.

This is the comparison a sponsor rarely sees laid out. A lawyer asked to solve the exempt investors' problem will reach for one of these two, and either will work. The value of the exploration is in showing what the choice costs, who bears it and what would have to be true for the other candidate to be better.

5.3 Adding non-U.S. investors: the diagrams hold, the space moves

Now add non-U.S. investors. The pass-through portfolio generates income effectively connected with a U.S. trade or business, which exposes those investors to U.S. tax and filing obligations and the fund to withholding on their share [24]. The single partnership now carries a third finding. The two candidates from the previous stage remain the credible responses, and their diagrams are unchanged: the non-U.S. investors simply join the exempt investors on the blocked path.

Their consequences are not unchanged. The modeled leakage attributed to the blocked sleeve rises, chiefly because more of the fund's capital now passes through the corporation rather than because the tax borne per dollar has changed, and distributions from a U.S. corporation to non-U.S. holders bring withholding considerations of their own. The sleeve now contains two classes of investor whose reasons for being there differ. A non-U.S. pension fund, a sovereign investor relying on its own exemption and a non-U.S. individual do not have identical requirements: they differ in treaty position, in willingness to accept U.S. filing obligations and in the rights they will negotiate. Whether they can share one vehicle depends on whether those differences can be accommodated without imposing a cost on the exempt investors already there. The blocked feeder as modeled, in which everyone shares one investment path, is under more strain from this addition than the parallel fund, even though both diagrams look the same as they did a stage earlier.

The offshore master-feeder that a sponsor might reach for on hearing "non-U.S. investors" is a feasible alternative but not a competitive one. An offshore vehicle does nothing about effectively connected income from U.S. operating businesses; the location of the feeder is not what determines the result, the participation path is. In the modeled scenario it carries materially higher leakage than either of the two candidates just discussed. A geographic

label is not a substitute for tracing the consequences through the arrangement.

5.4 Three patterns the exploration reveals

The example was chosen because a funds lawyer already knows that exempt and non-U.S. investors raise these issues. Its point is what exploring the design space adds once those issues are on the table, and three patterns stand out.

A structure can remain unchanged while its consequences change materially. Between the second and third stages the leading diagrams did not move, but the tax borne on the blocked path, the composition of the sleeve and the strain on a single set of terms all did. A review that asks only whether the architecture should change would report nothing; a review that tracks consequences reports a great deal. This is the point made abstractly in Section 2.3, and it is the reason a model should treat a fact as material when it changes findings or costs, not only when it changes the shape of the answer.

A favorable aggregate result can conceal an unfavorable outcome for one investor class. The fund-level leakage figures of the two candidates are similar and modest. All of that leakage is borne by the blocked sleeve. An exploration that reports only the fund-level number would not tell an exempt investor what its participation actually costs, or tell the sponsor that a term negotiated by the taxable investors is being paid for by someone else.

A structure suited to the expected launch conditions may be fragile under plausible changes. The choice between the parallel fund and the blocked feeder turns on a commitment threshold that itself depends on cost assumptions; the blocker is essential for a leveraged pass-through portfolio and becomes largely carrying cost if the portfolio shifts toward corporate holdings and the acquisition debt moves to the portfolio-company level; both additional vehicles are wasted expense if the exempt capital never arrives. Each of those is a plausible development for a first-time fund. A design that is stated together with the conditions under which it should be revisited is more useful than one stated as a conclusion.

6 Interpretability and professional review

A computational conclusion is useful to counsel when its basis can be examined. The model associates findings with the relevant parts of a design and provides a reasoning trace identifying the facts, characterized rules and authorities supporting each modeled conclusion, so

that a reviewer can isolate the source of a disagreement. Each fact in that trace should carry its basis, stated by the user or supplied by the model, so that the reviewer can see at once whether a disagreement is with an input or with an interpretation.

A reasoning trace states what supports a conclusion. A reviewer often wants the converse: what a particular vehicle in the structure is doing there. The direct way to answer is a counterfactual. Remove the vehicle from a copy of the design, re-run the characterized rules and report what appears, what disappears and what is only partly affected. That answer is computed rather than recalled, and it establishes something specific: that in the current structure the vehicle addresses an identified exposure. It does not establish why anyone put it there, nor that it is the best response available, and an explanation built on it should not claim either.

Language models have a place in this review as translators. A model can restate a finding, a counterfactual or a comparison between candidates in plain terms and connect it to the facts it rests on, provided that every factual claim in the explanation points back to a finding the rules produced. A model that is permitted to add findings in the course of explaining them has become the engine by another route, and the explanation has lost the property that made it reviewable.

The disagreement may concern an input, a legal interpretation or a commercial preference, and those are different problems. Correcting a factual assumption may change the analysis throughout the structure. Revising an interpretation may affect several otherwise unrelated scenarios. Choosing to accept an identified exposure may leave the legal analysis intact while changing the preferred commercial response.

Professional review includes confirming the facts, examining consequential interpretations and identifying requirements the model has not addressed. Counsel must also translate the chosen design into agreements and an operating arrangement that preserve its intended treatment. The model organizes the basis for that work and makes alternative choices easier to discuss.

7 Scope, validation and limits

The scope of a computational model has several dimensions: the facts it can represent, the rules it has characterized, the structures it can examine and the interactions it can evaluate. Broad coverage in one dimension does not establish completeness in the others.

A confirming finding means that a particular modeled condition is satisfied on the stated facts; it does not mean that every relevant issue has been examined, and an absence of warnings cannot distinguish a design with no identified concerns from one that presents issues outside the model's coverage. A model should therefore state its coverage so that a reviewer can tell what lies outside it.

Repeatability, legal correctness and commercial suitability are separate claims. Versioned inputs and rules support repeatability. Legal correctness requires review of the authority, the interpretation and its application to the facts; an explicit representation makes an interpretation inspectable but does not make it correct, and where the law depends on judgment or unresolved facts the conclusion should preserve that dependency. Commercial suitability depends on investor requirements and sponsor objectives that may extend beyond the model.

Where a language model assists the search or the explanation, its contribution should be labeled as such. A reviewer should be able to tell, for any statement in the output, whether it is a conclusion of the characterized rules on the stated facts, a hypothesis the model proposed and the rules tested, a matter for professional judgment or a question outside the model's coverage. The rules' answers should not depend on the model being available at all. A bounded search has a limit of its own: that it found nothing better than the generated candidates does not establish that nothing better exists, only that nothing better was found within the proposals it was allowed to make.

Validation of the implementation takes several forms. Reference structures that experienced counsel would recognize can be checked on every build. Regression tests can detect unintended changes in behavior, and doctrine tests can check that the implementation preserves specified relationships or conclusions. These methods test the implementation against intended behavior; they do not independently establish that the underlying legal interpretation is correct.

8 Implications for managers and counsel

Computational exploration of the design space can change how early a sponsor examines the consequences of a proposed design, and who takes part. An emerging manager can explore the implications of its expected investor base before committing to an architecture. An established manager can examine whether a new activity or investor requirement warrants revisiting an existing arrangement. A placement agent or product team can test

what the vehicle the market is asking for would cost and require before the legal team is engaged to build it, so that the demand side and the implementation side are working from the same set of alternatives rather than meeting after the architecture is fixed.

For counsel, a common representation of the design makes assumptions and disagreements easier to locate. Expert attention can focus on consequential facts, competing interpretations and the trade-offs among credible alternatives. The value of a mapped design space depends on preserving that visibility rather than allowing a condensed recommendation to obscure the reasons behind it.

A further implication follows from treating the structure as an explicit object. Once a design is represented as a graph carrying the facts, roles, jurisdictions, classifications and modeled conclusions that produced it, that representation can serve as the source of truth for everything that follows. Today it organizes the findings, the supporting authorities and the documents a structure requires. As language models become more capable of drafting, the same representation can supply the facts and design decisions the documents must implement, so that drafting begins from a confirmed architecture rather than from a precedent that may embed different assumptions. Whether that is realized depends on the reliability of the drafting tools and on professional review of their output. The design principle is the same in either case: the structure is decided first, made explicit and carried forward, and the documents follow from it.

9 Structura

The principles described in this paper are the ones on which we built Structura, a fund architecture platform developed by AnchorPoint Labs LLC. Sections 2 through 6 describe the framework; this section states what Structura implements today.

Structura represents a proposed fund as the graph described in Section 2, with typed investors, entities and holdings and their relationships drawn on a canvas, and records for each fact whether the user stated it or the model supplied it. It applies a versioned library of characterized legal rules to those facts, attaching each finding to the part of the structure it concerns together with its supporting authority, the facts it relied on where the rule records them, and an indication of whether the rule rests on settled authority, an interpretation or authority that is still moving. It generates the candidate structures within its coverage for the stated facts, the

modeled design space of Section 3, and compares them, associating formation and recurring cost assumptions and a modeled economic outcome with each candidate and reporting the modeled tax leakage by participation path rather than only in aggregate. It re-runs the exploration with one modeled fact changed at a time and reports the conditions under which the preferred candidate changes, which is the analysis behind Section 5.4. It provides a reasoning trace for each candidate identifying the facts, rules and authorities supporting each modeled conclusion, answers what a vehicle is doing in a structure by the counterfactual of Section 6, and derives the documents the chosen structure requires. Figure 1 shows a candidate structure as drawn on the canvas.




Figure 1. A candidate structure for the illustrative buyout fund of Section 5, as drawn on the Structura canvas: a main fund and a parallel fund under a common general partner and management company, with the tax-exempt investors' participation routed through a corporate blocker.

Language models are used in the four roles described in Sections 2 and 6 and in no other. They assist in developing the rule library. At intake they translate a manager's description, or an existing structure chart, into facts that the user confirms, asking for what is missing rather than supplying it. In the search, a model proposes bounded changes to a structure from a fixed vocabulary of operations, such as interposing a blocker, adding a feeder or a parallel vehicle, changing a jurisdiction or a tax election, or removing a vehicle; Structura applies each proposal to a copy of the structure, runs the rules, the cost model and the economic model on the result and ranks it among the other candidates on the same basis. The deterministic ranking is produced first and stands on its own; the search is an optional second phase of the same

analysis, and a variant it finds is shown as a proposal the model made and a result Structura computed. In explanation, a model restates findings, counterfactuals and candidate comparisons in plain terms, with every factual claim tied to a finding the rules produced. The model does not change a fact, score a candidate, state a rule or create a finding. Its observations about a structure are labeled as observations and leave the findings unchanged, and every conclusion Structura reports is available whether or not a model was called.

As of October 2026, the model covers eight investor tax classifications and forty-two investment activities across twelve asset groups, from buyout and venture through private credit, real estate, infrastructure, secondaries and liquid strategies. The rule library addresses securities offerings, investment company and adviser regulation, federal tax and ERISA, together with a state layer covering adviser registration, offering filings and partnership tax requirements in every U.S. jurisdiction. Every rule carries its authority, its review date and its classification as settled or interpretive, and the library is versioned so that any output can be reproduced against the rules that produced it. The rule library and the range of structures the model can generate grow with each release.

Notes

1. The earliest system in artificial intelligence and law was already about tax structures: TAXMAN classified corporate reorganizations under the Internal Revenue Code, taking a transaction as given and analyzing it [25]. Planning followed. Michaelsen's TAXADVISOR encoded federal income and transfer tax planning knowledge as an expert system and was tested against practitioners [26]; EPS II and McCarty's later statement of the general form are cited in Section 3 [17, 18, 19]. Sanders' CHIRON generated income tax plans from statutes and cases in an open-textured domain [27]. The field's own retrospective describes legal planning as a trail those papers blazed and later work did not follow [28], and the most recent legal-planning work, on autonomous vehicles and the highway code, still finds the problem little studied [29]. Several recent lines of tax work run parallel: search over partnership transaction structures against a simulated tax code to find abusive schemes [30]; network analysis of the treaty system to find the lowest-tax route for repatriated profits [31]; an automated planner over a small vocabulary of entities, ownership and licensing that reconstructs a well-known international avoidance structure [32]; and, closest in its division of labor, a 2026 system in which a language model translates the tax problem and a constraint solver searches for a compliant optimum, although what it optimizes is tax variables under a statute rather than which entities to create and how investors participate through them [33]. The law-review

literature has, in parallel, examined how the Code itself can be formalized and what formalization makes possible [34, 35, 36].

References

- [1] C. Mead and L. Conway, *Introduction to VLSI Systems* (Addison-Wesley, 1980).
- [2] H. A. Simon, *The Sciences of the Artificial*, 3rd ed. (MIT Press, 1996).
- [3] D. Cassens Weiss, "This BigLaw Firm Charges Nearly \$2,500 per Hour for Top Billers' Bankruptcy Work," *ABA Journal* (Dec. 19, 2023), reporting Kirkland & Ellis rate disclosures in court filings: partners \$1,195 to \$2,465 and associates \$745 to \$1,495 per hour.
- [4] D. Cassens Weiss, "Lawyer Fees Skyrocket; Some Charge up to \$3,400 per Hour," *ABA Journal* (Feb. 19, 2026), reporting Wall Street Journal and Persuit data on partner rates at the fifty largest U.S. law firms.
- [5] Institutional Limited Partners Association, *ILPA Principles 3.0: Fostering Transparency, Governance and Alignment of Interests for General and Limited Partners* (June 2019).
- [6] Institutional Limited Partners Association, *The Alignment Gap: Rethinking Costs in Private Equity Fund Formation* (May 14, 2026), <https://ilpa.org/industry-guidance/principles-best-practices-organizational-expenses/>.
- [7] J. S. Levin and D. E. Rocap, *Structuring Venture Capital, Private Equity and Entrepreneurial Transactions* (Wolters Kluwer, updated annually).
- [8] M. J. Sergot, F. Sadri, R. A. Kowalski, F. Kriwaczek, P. Hammond and H. T. Cory, "The British Nationality Act as a Logic Program," 29 *Communications of the ACM* 370 (1986).
- [9] M. Genesereth, "Computational Law: The Cop in the Backseat," *Stanford Center for Legal Informatics (CodeX)* (2015).
- [10] OECD, *Cracking the Code: Rulemaking for Humans and Machines*, OECD Working Papers on Public Governance (2020).
- [11] D. Merigoux, N. Chataing and J. Protzenko, "Catala: A Programming Language for the Law," 5 *Proceedings of the ACM on Programming Languages (ICFP)* art. 77 (2021).
- [12] *Soroban Capital Partners LP v. Commissioner*, No. 25-2079 (2d Cir. Sept. 17, 2026), affirming 161 T.C. 310 (2023) and T.C. Memo. 2025-52, <https://law.justia.com/cases/federal/appellate-courts/ca2/25-2079/25-2079-2026-09-17.html>.
- [13] T. Bench-Capon and F. Coenen, "Isomorphism and Legal Knowledge Based Systems," 1 *Artificial Intelligence and Law* 65 (1992).
- [14] M. Dahl, V. Magesh, M. Suzgun and D. E. Ho, "Large Legal Fictions: Profiling Legal Hallucinations in Large Language Models," 16 *Journal of Legal Analysis* 64 (2024).
- [15] V. Magesh, F. Surani, M. Dahl, M. Suzgun, C. D. Manning and D. E. Ho, "Hallucination-Free? Assessing the Reliability of Leading AI Legal Research Tools," 22 *Journal of Empirical Legal Studies* 216 (2025).
- [16] 15 U.S.C. § 80b-3(l) (exemption for advisers solely to venture capital funds) and 17 C.F.R. § 275.203(l)-1 (definition of venture capital fund; qualifying investments and the twenty percent non-qualifying basket).

- [17] D. A. Schlobohm and L. T. McCarty, “EPS II: Estate Planning with Prototypes,” *Proceedings of the Second International Conference on Artificial Intelligence and Law* (ACM, 1989), 1.
- [18] D. A. Schlobohm and D. A. Waterman, “Explanation for an Expert System That Performs Estate Planning,” *Proceedings of the First International Conference on Artificial Intelligence and Law* (ACM, 1987), 18.
- [19] L. T. McCarty, “Some Requirements on an Action Language for Legal Discourse,” *AAAI Spring Symposium Series, Technical Report SS-95-07* (AAAI Press, 1995).
- [20] R. L. Keeney and H. Raiffa, *Decisions with Multiple Objectives: Preferences and Value Tradeoffs* (Wiley, 1976; Cambridge University Press, 1993).
- [21] A. Saltelli, M. Ratto, T. Andres, F. Campolongo, J. Cariboni, D. Gatelli, M. Saisana and S. Tarantola, *Global Sensitivity Analysis: The Primer* (Wiley, 2008).
- [22] R. de Neufville and S. Scholtes, *Flexibility in Engineering Design* (MIT Press, 2011).
- [23] 26 U.S.C. §§ 511 to 514 (unrelated business taxable income; unrelated debt-financed income).
- [24] 26 U.S.C. §§ 864(b), 871(b), 875, 882 and 1446 (income effectively connected with a U.S. trade or business; partners in partnerships so engaged; withholding on foreign partners’ distributive shares).
- [25] L. T. McCarty, “Reflections on TAXMAN: An Experiment in Artificial Intelligence and Legal Reasoning,” 90 *Harvard Law Review* 837 (1977).
- [26] R. H. Michaelsen, “An Expert System for Federal Tax Planning,” 1 *Expert Systems* 149 (1984).
- [27] K. E. Sanders, “CHIRON: Planning in an Open-Textured Domain,” 9 *Artificial Intelligence and Law* 225 (2001).
- [28] T. Bench-Capon et al., “A History of AI and Law in 50 Papers: 25 Years of the International Conference on AI and Law,” 20 *Artificial Intelligence and Law* 215 (2012), commentary on EPS II by K. D. Ashley.
- [29] G. Sartor, G. Governatori, G. Pisano, A. Rotolo and A. Wyner, “Plans and Diversions,” in *Legal Knowledge and Information Systems: JURIX 2025, Frontiers in Artificial Intelligence and Applications* (IOS Press, 2025).
- [30] E. Hemberg, J. Rosen, G. Warner, S. Wijesinghe and U.-M. O’Reilly, “Detecting Tax Evasion: A Co-evolutionary Approach,” 24 *Artificial Intelligence and Law* 149 (2016).
- [31] M. van ’t Riet and A. Lejour, “Optimal Tax Routing: Network Analysis of FDI Diversion,” 25 *International Tax and Public Finance* 1321 (2018).
- [32] P. Fratrič, N. Holzenberger and D. Restrepo Amariles, “Can AI Expose Tax Loopholes? Towards a New Generation of Legal Policy Assistants,” arXiv:2503.17339 (2025).
- [33] T.-C. Hwang, F. Yu and J.-H. R. Jiang, “Agentic Taxation Optimization via LLM SMT-Constraint Reasoning,” *Proceedings of the IEEE/ACM 48th International Conference on Software Engineering: Software Engineering in Practice* (ACM, 2026).
- [34] S. B. Lawsky, “Formalizing the Code,” 70 *Tax Law Review* 377 (2017).

[35] S. B. Lawsky, “A Logic for Statutes,” 21 *Florida Tax Review* 60 (2017).

[36] S. B. Lawsky, “Tax Law and Flexible Formalizations,” *University of Chicago Law Review Online* (Nov. 18, 2024).

About this paper

Structura was built by AnchorPoint Labs LLC, an affiliate of AnchorPoint Legal Partners, PLLC, a law firm focused on private funds and recognized by Chambers and Partners.

This paper is for general information only. It is not legal, tax or investment advice and does not create an attorney-client relationship. A proposed structure should be reviewed by qualified counsel on the facts of the particular fund. Model coverage and outputs may change.