

Doctoral Advanced Module — Partnerships and Ecosystems as a Research Domain

The federal market is not won by firms alone. It is won by teams, joint ventures, mentor-protégé pairs, and whole ecosystems — and the discipline treats partnering mostly as compliance and occasionally as strategy. This seminar treats it as a research domain: why alliances form, how workshare works, whether the ecosystem compounds, and whether partnerships actually perform after the award.

Level: doctoral (L3 Architect / L4 Networker) · **Format:** one 3-hour seminar per week, five sessions · **Prerequisites:** DL 901 (design and measurement) · **Companions:** doctrine/06, doctrine/07, course/graduate-elective-teaming-jv-subcontracting.md, modules/doctoral/research-agenda.md · **Stack-agnostic promise:** partnering and ecosystem structure are studied as durable relationships and structures — never as a product’s workflow (doctrine/08).

This module is the third of the three C.A.S.E.-adjacent advanced modules the degree crosswalk promises at the doctoral tier (alongside AI systems and advanced capture management). Where the master’s elective teaches the *craft* of teaming — the partner rubric, the agreements, the gap analysis — this seminar studies partnering and ecosystems as *research objects*: the theory of why alliances form, the dynamics of workshare, the ecosystem as a unit of analysis, and the post-award performance of partnerships.

Course description

Partnerships are the discipline’s most under-theorized leverage. Every serious pursuit is a teaming decision: prime or sub, JV or mentor-protégé, which partners, what workshare, what structure. The decision is made under a heavy regulatory structure (small-business programs, size standards, JV rules) and a thin evidence base — firms argue about partnering with anecdotes, and the field publishes almost nothing systematic on whether partnering actually works. This seminar converts the argument into research.

The seminar moves through four questions:

1. **Why do alliances form?** Alliance formation theory — transaction-cost, resource-based, and network explanations — gives the discipline a theory of when partnering is strategic rather than compliance.
2. **How does workshare work?** The prime/sub/JV structure and the workshare allocation are a design decision with measurable consequences.
3. **Is the ecosystem real?** The Dream ecosystem — a platform that turns dreams into orbitals into worlds — is a researchable case of ecosystem structure, with a canary company as a shared test object.
4. **Do partnerships perform?** Post-award, does teaming, JV, or mentor-protégé structure predict delivery outcomes — and does past performance

compound through partnerships?

The deliverable is a **seminar research design** — one study, fully designed, defensible to a skeptical panel — that sits on the research agenda and is designed to grow into the dissertation’s empirical core.

Learning outcomes

By the end of this module, a student can:

1. **Explain** alliance formation theory — transaction-cost, resource-based, and network accounts — and apply it to the discipline’s teaming decisions.
2. **Operationalize** workshare as a research construct — the allocation of work, value, and risk across partners — and design the study of its consequences.
3. **Treat** an ecosystem (the Dream ecosystem as a case) as a unit of analysis, with position, platform structure, and the cosmology loop as measurable features.
4. **Design** a post-award performance study of partnership structures, connecting delivery outcomes to past-performance compounding.
5. **Defend** a seminar research design — including the partnership-strategy deliverable the competency model calls for at the doctoral tier — before a skeptical panel.

Session 1 — Alliance formation theory: why firms partner

The claim. Teaming is not just a compliance exercise — it is a strategic choice with a theory. The alliance literature explains why firms form alliances, when they do, and why some alliances create value while others are theater.

The ideas. Three scholarly accounts, applied to the federal market:

- **Transaction-cost economics.** A firm partners when the cost of organizing a capability through the market (finding, contracting, governing a supplier) is higher than the cost of organizing it through an alliance. In the federal market, where size standards and set-asides make certain work *legally unavailable* to a firm alone, the transaction-cost account is unusually concrete: the small business that cannot bid without a mentor or a prime is partnering because the market for that work is closed to it.
- **The resource-based view.** A firm partners to access capabilities it does not own — past performance, an incumbent relationship, a facility clearance, a named anchor partner. The alliance is a vehicle for resource access, and the partner’s value is the resource it brings.
- **Network position.** A firm’s position in a network of relationships — primes, subs, agencies, research institutions — is itself an asset. Partnering

is how a firm moves through the network; the network is how the market is actually structured.

The hidden assumptions. That the alliance is formed for value rather than compliance; that the partner brings the resource it claims; that the structure (teaming agreement, JV, mentor-protégé) matches the strategic intent; that trust can substitute for formal governance — and vice versa.

The research questions.

- When is teaming **strategic** and when is it **compliance** — and can the two be distinguished in the record?
- What predicts **alliance success** in federal competitions — resource complementarity, prior relationship, workshare design, governance structure?
- Does **ecosystem position** — the firm's place in the web of primes, subs, and agencies — predict win rates, controlling for firm quality?

Design exercise. Teams design the **alliance-formation study**: the population (firm-competition pairs with teaming configurations, drawn from the public award record and contractor registrations), the measures (partner complementarity, prior relationship, structure type), the outcome (win/loss, and post-award delivery quality where observable), and the confounders (large firms partner more; quality firms attract partners).

Reading. doctrine/01 (the small-business and set-aside structure — where partnering is legally required); course/graduate-elective-teaming-jv-subcontracting.md at altitude (the craft the research explains); the alliance-theory canon as named concepts — transaction-cost economics, the resource-based view, network theory.

Session 2 — Workshare dynamics: the structure of the alliance

The claim. The workshare allocation — who does what work, who holds what value, who carries what risk — is the alliance's design decision, and it has measurable consequences for pursuit quality and delivery performance.

The ideas. A partnership structure is a bundle of design choices:

- **Structure type** — prime/sub (one firm holds the prime contract, the other performs), joint venture (a separate legal entity, each partner holds a share), mentor-protégé (a small-business program pairing, with developmental commitments beyond the single pursuit). Each type allocates work, value, and risk differently.
- **Workshare allocation** — the split of the work and the value across partners. The discipline teaches that workshare should follow *who brings the discriminating capability*, not *who has the larger back-office* — a claim that is directly testable.

- **The governance of the alliance** — the teaming agreement, the decision rights, the dispute path. Governance structure is the alliance’s gate structure (doctrine/04).

The hidden assumptions. That workshare follows capability rather than politics; that the JV is a real operating entity rather than a compliance shell; that the mentor actually develops the protégé; that the alliance’s governance survives the first hard decision.

The research questions.

- Does **workshare design** predict pursuit outcomes — do proposals where the discriminating capability owns the largest share win more?
- Does **structure type** matter — do JVs, prime/sub teams, and mentor-protégé pairs perform differently, controlling for the underlying firms?
- Do **mentor-protégé pairs** compound — does the protégé’s win rate and capability rise over the program, and does the mentor actually develop or merely sponsor?

Design exercise. Teams design the **workshare-design study**: the data (teaming configurations and workshare splits from proposals and award records, where disclosed, plus organizational records where obtainable), the measures (workshare concentration, capability-share alignment, structure type), the outcome (win/loss and delivery), and the selection hazard (firms select into structures non-randomly).

Reading. doctrine/04 (gates as the governance frame) and doctrine/06 (structure-first — the ORBITAL and its mapping to proposal sections); the JV and mentor-protégé program concepts (13 CFR small-business joint ventures, the SBA mentor-protégé program) as public-rule references; course/graduate-elective-teaming-jv-subcontracting.md.

Session 3 — The ecosystem as a case: structure, position, and the cosmology loop

The claim. The Dream ecosystem — a multi-organization platform that turns dreams into orbitals into worlds — is a researchable case of ecosystem structure. Its claims — that ecosystem position compounds, that the cosmology loop re-seeds discovery — are hypotheses, not assertions.

The ideas. An ecosystem is a network of organizations that co-specialize: a platform, its complementors, and the shared infrastructure that lets them create value together. The Dream ecosystem is a deliberately designed instance: a pipeline (sense → qualify → score → pursue → design → audit → submit → deliver → learn), a shared doctrine encoded in two renderings, a canary company (Ravonics) used to validate the whole system, and a cosmology loop that promises a Won-World re-seeds new Dreams (doctrine/07). For a researcher, this is a *case* — a bounded, designed system whose claims can be studied:

- **The platform claim.** That the shared doctrine and pipeline raise the win rate of every organization on it, not just the operator.
- **The position claim.** That where an organization sits in the ecosystem — its relationships, its past-performance record, its place in the pipeline — predicts its next probability of win.
- **The cosmology claim.** That won work compounds — that the ecosystem’s outcomes feed back and raise the next generation’s odds (Thread E).

The hidden assumptions. That the ecosystem is more than the sum of its members; that position is measurable and not just a story; that the cosmology loop actually fires; that the canary company’s success generalizes.

The research questions.

- Does **ecosystem position** predict win rates — do organizations with richer network position, past-performance records, and pipeline standing win more, controlling for firm quality?
- Does the **cosmology loop** compound — does won work measurably raise the next probability of win, in the way the doctrine claims (Thread E)?
- Is the **canary-company pattern** a valid evaluation device — does validating the pipeline on one canonical company prove anything about the world (the methodology thread)?

Design exercise. Teams design the **ecosystem-position study**: the data (ecosystem membership, relationship graphs, past-performance records, pursuit outcomes over time), the design (a panel study of ecosystem members’ win rates as their position changes, or a comparative case of the Dream ecosystem against a non-ecosystem counterfactual), and the survivorship and selection hazards.

Reading. doctrine/07 (the lifecycle and the cosmology loop); case-study/ravonics.md (the canary company as the shared test object); align/concept-to-system-map.md (the ecosystem’s own architecture, at concept level); research-agenda Threads D and E.

Session 4 — Post-award performance of partnerships

The claim. Partnerships are not just pursuit structures — they are delivery structures. The question the field rarely asks is whether teaming, JV, or mentor-protégé structure predicts what happens *after* the award: performance, past performance, and the compounding that feeds the next pursuit.

The ideas. Post-award performance is the field’s most under-studied outcome because it is harder to observe than a win/loss. But it is researchable: past performance records, delivery outcomes, contract outcomes (option exercise, follow-on awards, protests), and — for JVs and mentor-protégé pairs — the trajectory of the partner firms themselves. The theoretical frame connects to Thread E (learning and evolution): the discipline claims that past performance is

the compounding input (doctrine/10), and partnerships are the vehicle through which a small firm gains the past performance it cannot build alone.

The hidden assumptions. That the partnership that wins is the partnership that performs; that a JV that wins actually delivers rather than dissolving after award; that the mentor-protégé pair's developmental gains are real; that post-award outcomes can be attributed to structure rather than to the firms' underlying quality.

The research questions.

- Does **partnership structure** predict delivery outcomes — do JVs, prime/sub teams, and solo primes differ on past performance, option exercise, and follow-on awards?
- Does the **winning partnership perform** — or do partnerships dissolve after award, and does dissolution predict delivery problems?
- Does **mentor-protégé** participation compound — do protégé firms show measurable capability and win-rate gains over the program, and do those gains survive the program's end?
- Does **partnered past performance** transfer — when a small firm wins through a prime, does the resulting past performance record improve its next solo probability of win?

Design exercise. Teams design the **post-award partnership-performance study**: the data (award records joined to past performance and delivery outcomes — public where available, organizational where not), the design (a cohort or panel study of partnership structures and their delivery trajectories), and the attribution hazard (firms select into structures; quality confounds).

Reading. doctrine/10 (post-submission and win — past performance as the compounding input); doctrine/07 (the lifecycle — delivery as a stage); the past-performance retrieval concepts from align/concept-to-system-map.md §7; research-agenda Thread E.

Session 5 — Designing the partnership study: synthesis and defense

The final session is a working session and a defense. The four questions — why alliances form, how workshare works, whether the ecosystem compounds, whether partnerships perform — are one story: the alliance is formed for a reason (Session 1), structured by a workshare design (Session 2), embedded in an ecosystem (Session 3), and judged by its delivery (Session 4). Each student presents their seminar research design — the study that tests one claim about partnerships or ecosystems — and the room attacks it the way a dissertation committee would: *What is the selection hazard? Is the sample big enough? How do you measure workshare without access to internal agreements? What exactly will you claim if the data go the other way?*

The panel's job is the committee's job: push until the design's limits are named, then judge whether the design is *defensible*, not perfect. For students pursuing the partnership-strategy track, this session also walks the **partnership-strategy deliverable** the competency model names at the doctoral tier: a teaming/posture framework for an organization across a portfolio — prime/sub/JV/mentor-protégé policy, partner rubric, workshare design — argued against the evidence the seminar has read.

Research-question bank

Each question is dissertation-sized: researchable, feasible, and a contribution the field needs.

1. **Strategic vs. compliance teaming.** When is teaming a strategic choice and when is it a compliance requirement — and can the two be distinguished in the record?
2. **Alliance success.** What predicts alliance success in federal competitions — resource complementarity, prior relationship, workshare design, or governance structure?
3. **Workshare and win.** Does the workshare allocation — whether the discriminating capability owns the largest share — predict pursuit outcomes?
4. **Structure type.** Do JVs, prime/sub teams, and mentor-protégé pairs perform differently, controlling for the underlying firms?
5. **Ecosystem position.** Does an organization's position in an ecosystem — its relationships, past-performance record, and pipeline standing — predict its win rate?
6. **The cosmology loop.** Does won work compound — does ecosystem success measurably raise the next generation's probability of win?
7. **The canary pattern.** Is validating a pipeline on one canonical company a valid evaluation device, and what does a green run certify?
8. **Post-award structure.** Does partnership structure predict delivery outcomes — past performance, option exercise, follow-on awards — or does the underlying firm quality explain it all?
9. **Partnership dissolution.** Do partnerships dissolve after award, and does dissolution predict delivery problems?
10. **Mentor-protégé compounding.** Do protégé firms show measurable capability and win-rate gains over the program — and do the gains survive the program's end?
11. **Past-performance transfer.** When a small firm wins through a prime, does the resulting past performance improve its next solo probability of win?

Reading list

The module is concept-first; the reading is named concepts and schools, not pirated material:

- **The doctrine as the discipline's claims:** doctrine/01 (the small-business and set-aside structure), doctrine/04 (gates as governance), doctrine/06 (structure-first — the ORBITAL), doctrine/07 (the lifecycle and the cosmology loop), doctrine/10 (past performance as the compounding input).
- **The alliance and network canon:** transaction-cost economics (the make-vs-buy account of when firms partner), the resource-based view (partnership as resource access), network theory (position as an asset), and the inter-organizational-alliance literature as named concepts.
- **The partnering craft at altitude:** course/graduate-elective-teaming-jv-subcontracting.md (the partner rubric, agreements, gap analysis — the craft the research explains); the JV and mentor-protégé program concepts from the public rules (13 CFR small-business joint ventures; the SBA mentor-protégé program).
- **The ecosystem frame:** the platform-and-ecosystem concepts — co-specialization, complementors, network effects — as named schools; and the Dream ecosystem's own architecture as rendered in align/concept-to-system-map.md and case-study/ravonics.md.
- **The methods core:** DL 901 (design and measurement — including comparative case and network designs) and DL 902 (statistics and causal inference) — the toolkit this module points at the alliance question.
- **The public record as evidence:** the award record, contractor registrations, and the past-performance record — the ground truth for partnership research.

Dissertation tie-in

This module is the natural home of dissertations on the partnership-strategy deliverable the competency model names at the doctoral tier (design a teaming/posture framework for an organization across a portfolio), with direct bridges into Thread E (learning and evolution — the cosmology loop) and Thread D (decision and governance — the alliance as a governed structure). A dissertation from this module typically lands on one: an alliance-formation study (when is teaming strategic?), a workshare-design study (does workshare predict wins?), an ecosystem-position study (does position compound?), or a post-award performance study (does partnering deliver?). The session designs assemble directly into a dissertation proposal.

The seminar in two sentences

Partnerships are the discipline's most under-theorized leverage — every pursuit is a teaming decision, and the field argues about teaming without evidence. This

seminar trains the researchers who will close that gap: the alliance theorized, the workshare measured, the ecosystem tested, and the partnership judged by its delivery. **Partnering is a claim about leverage — and every claim can be tested.**

Reflection questions

1. Which claim about partnerships do you most suspect is *not* true — that workshare follows capability, that mentor-protégé compounds, that ecosystem position predicts wins — and what would your study show if your suspicion were right?
2. For the design you are building, what is the single most likely way a skeptic could dismiss the result, and what have you done to pre-empt it?
3. If a BD executive reads your published study, what is the one sentence of guidance about partnering they should be able to take away — and is your design strong enough to support it?