

Post-Doctoral Fellowship — The Research Career After the PhD

The doctorate proved you can test the doctrine. The fellowship is where you take responsibility for it — owning the research agenda, authoring and curating the corpus, and raising the next generation of researchers and practitioners.

Level: post-doctoral · **Companions:** [course/doctoral-program-overview.md](#), [modules/doctoral/research-agenda.md](#), [course/forever-learning-program.md](#).

This module is the research career after the doctorate. It is written for the person who finished a dissertation and did not stop — the early-career researcher who wants to keep publishing, the senior practitioner with a research arm, the faculty member building a program, or the retired capture professional who wants to leave the discipline better than they found it. It tells you what the fellowship is, what it expects of you, and how the discipline makes room for a lifetime of research.

1. What a post-doctoral fellowship is here

In a university, a post-doctoral fellowship is a temporary research appointment after the PhD — a chance to publish, to learn new methods, and to position for a faculty role. This tier keeps that shape but **deliberately does not require a university**. A post-doctoral fellowship in the Dream pursuit doctrine is a *standing research commitment* with four privileges and four responsibilities:

The privileges	The responsibilities
Access to the research agenda and the outcome-data record	Own at least one research thread (see §4)
A peer seminar of fellow researchers	Publish on a cadence — and defend the discipline’s integrity standards
Faculty standing in the program (teach, advise, chair)	Mentor doctoral candidates through their milestone gates
Editorial voice over the DNA corpus	Author and curate corpus doctrine — and keep it versioned and honest

The fellowship is the “Global Impact” row of the doctoral overview’s ladder made real: where the doctoral student *tests* the doctrine, the fellow *stewards* it — decides which questions matter, shapes the corpus, and trains the people who will run the field.

2. The advanced research seminar

The heart of the fellowship is the **advanced research seminar** — a recurring, peer-reviewed working group of post-doctoral fellows, doctoral candidates, and

invited senior practitioners. Unlike the doctoral coursework (DL 901–903), the seminar is not graded and has no fixed syllabus. It is a **standing structure** with three standing items:

1. **Work-in-progress review.** Each session, one fellow presents a study in progress — design, preliminary results, or a stalled analysis. The seminar’s job is the same as a color team’s (literacy/how-color-teams-work.md): independent readers, specific findings, a forward-only cadence. A research question gets the same ruthless peer treatment a proposal does.
2. **The machinery as research object.** The seminar keeps a running list of the discipline’s empirical claims and designs studies to test them. This is the DL 903 posture, continued for a career: every formula, every threshold, every gate rule in the doctrine is a claim that can be tested against the outcome record.
3. **Cross-disciplinary method swaps.** The seminar deliberately imports methods from neighboring fields — decision science, behavioral economics, game theory, operations research, organizational learning — and asks what each would say about a doctrine claim. This is how the discipline stays empirically young in the good sense: it keeps borrowing sharper tools.

The seminar’s output is not a grade; it is a **publication pipeline** — studies designed, executed, and submitted to peer-reviewed venues, plus the field’s shared methodological notes.

The year’s seminar arc

The advanced research seminar meets on a standing weekly rhythm, but its content follows a **year-long arc** — four blocks, each with its own objective, each training a distinct muscle the fellow will need. The arc keeps the fellowship’s research, corpus, and faculty work in one conversation, and it ensures that a fellow who joins mid-year can see where the seminar is in its rotation.

Block 1 — Methods (autumn, ~10 weeks). *The objective: sharpen the field’s empirical toolkit.* Each session imports or deepens one method — calibration measurement (reliability diagrams, Brier score, expected calibration error), causal inference (difference-in-differences, regression discontinuity), content analysis and inter-rater reliability, network methods, qualitative comparative case method — and applies it to a doctrine claim. The block runs on a method-swap format: a fellow or invited methodologist presents the method, and the seminar designs the study that would use it against a live question on the agenda. The block’s output is a running methodological note — the field’s shared toolkit, built in public.

Block 2 — Corpus authorship (winter, ~8 weeks). *The objective: turn tested claims into authored and curated doctrine.* The seminar shifts to the editorial bridge (see §3): a fellow presents a claim they have tested — a factor weight that does not survive across agencies, a gate rule that does not predict — and the seminar works it into a corpus proposal: the concept first (what should

be true), then the machine-side encoding (how a rule, formula, or playbook should express it). The block is also where curation is exercised: proposed machine-side changes are reviewed for doctrinal consistency, and the human-side sign-off on version bumps is argued. The block’s output is at least one corpus proposal per fellow, taken to a reviewable draft.

Block 3 — Supervision (spring, ~8 weeks). *The objective: practice the faculty and mentor track on live student work.* The seminar turns to the doctoral pipeline: a fellow presents the dissertation-stage work of a candidate they advise (or serves as the outside reader on a committee outside their thread), and the seminar rehearses the supervisor’s questions — *what is the identifying assumption, where is the survivorship hazard, is the sample big enough, what will you claim if the data go the other way?* The block trains the discipline of advising: reading a draft the way a chair reads it, holding a candidate to written criteria, and catching the expectation mismatch before the defense rather than at it.

Block 4 — Cross-disciplinary (late spring, ~6 weeks). *The objective: import sharper tools from the neighboring fields.* The seminar opens its doors to the fields the doctoral tier already teaches at altitude — decision science, behavioral economics, game theory, organizational learning — with invited readers or readings, and asks what each would say about a doctrine claim. The cross-disciplinary block is the fellowship’s standing commitment to staying empirically young in the good sense: it keeps borrowing tools (see §6), and it ends the seminar year by returning to the research agenda with new methods on the table.

The arc’s standing rhythm. Work-in-progress review runs *through* all four blocks — a session per month is reserved for a fellow presenting a study in progress, whatever the block’s theme. The arc organizes the seminar’s attention; the work-in-progress review keeps the publication pipeline moving all year. A fellow who completes the year has worked all four muscles — methods, corpus, supervision, cross-disciplinary — and has a publication pipeline, a corpus proposal, and an advising practice to show for it.

3. Contributing to the DNA corpus: authorship and curation

The discipline is built on “one doctrine, two renderings”: a human rendering (this curriculum) and a machine rendering (the DNA corpus — rules, formulas, playbooks, prompts, RAG material, versioned at `DreamLimited/dna`). The corpus does not write itself, and it must not drift from the human doctrine. The post-doctoral fellow is the discipline’s **editorial bridge** between the two renderings.

Two distinct roles:

- **Corpus authorship.** A fellow who has tested a claim — say, that a particular factor weight in the pWin composite does not survive across

agencies — is the person who should propose the improved rule. Authorship means writing the *concept* first (what should be true) and then working with the machine side to encode it (how a rule, formula, or playbook should express it). The human side is always the anchor; the machine side is always versioned.

- **Corpus curation.** More common than authorship is curation: keeping the corpus honest and current. A curator reviews proposed changes for doctrinal consistency, checks that a machine-side update does not silently change a durable concept, and signs off on version bumps from the human rendering’s side. Curation is a **gate-keeping role**, and it carries the discipline’s integrity standard: documented fact is separated from clearly-labeled illustrative material, and nothing is ever attributed to a real organization that did not do it.

The fellowship’s editorial voice is earned, not granted. A fellow begins by reviewing and curating, earns authorship on a tested claim, and only then holds the pen on a whole doctrine chapter.

4. The research agenda: owning a thread

The ORBITAL Outcome-Data Research Agenda organizes the field’s open questions into five threads. A doctoral dissertation is a deep contribution to **one** thread. A post-doctoral fellowship is a *commitment to a thread over years*:

Thread	What the fellow stewards
A — Prediction and calibration	The body of evidence on whether pWin estimates predict awards, and whether they are calibrated.
B — Factor validation	Which factors compose a win-probability estimate, and whether their weights hold across agencies and funder families.
C — Agency win patterns	Whether agencies’ published review criteria predict what they fund.
D — Decision and governance	Whether honest bid/no-bid gating, threshold discipline, and portfolio allocation improve outcomes.
E — Learning and evolution	Whether organizations that measure outcomes and learn from losses actually improve over time.

A thread owner does three things: keeps the thread’s literature map current, designs and executes the studies the thread needs, and recruits the doctoral candidates whose dissertations advance it. The fellowship is how the agenda stops being a document and becomes a **research program**.

5. The faculty and mentor track

The fellowship carries faculty standing in the curriculum. It is the teaching counterpart of the research role, and it uses the discipline's own governance ideas: milestones with written criteria, named gate owners, and a forward-only cadence.

- **Teaching.** A fellow teaches at the doctoral altitude — DL 901–903 — and at the forever-learning altitude (the advanced research seminar, executive master classes on research-informed practice). Teaching the doctrine is the surest test of understanding it: if you cannot teach a concept at concept altitude, you do not yet understand it.
- **Advising doctoral candidates.** The dissertation is a sequence of gates (course/doctoral-dissertation-sequence.md), and each gate has an owner. A fellow serves as a **committee member**, then as a **methodology advisor**, then as a **dissertation chair**. The progression is itself a gate discipline: a fellow does not chair until they have read enough drafts to know what a defensible contribution looks like.
- **Supervising research apprentices.** Before candidacy, doctoral students complete a research apprenticeship. The fellow is the supervisor — the person who gives a student a live question, holds the analysis to the discipline's standards, and models the non-fabrication and replication benchmarks the doctorate demands.
- **The external-reader role.** A fellow also serves as the outside reader on committees outside their thread — the insurance against insular thinking the doctoral overview calls for. A calibration researcher reading a governance dissertation is exactly the cross-fertilization the field needs.

6. Cross-disciplinary research: pWin, game theory, decision science

The discipline's open questions are not owned by capture management alone. The fellowship exists to keep importing sharper thinking from the fields the master's tier already teaches at altitude (course/graduate-competitive-strategy-game-theory.md, course/graduate-microeconomics-of-procurement.md):

- **Decision science.** The pWin estimate is, at bottom, a *judgment under uncertainty*. Decision science asks how people form such judgments and where they systematically over- or under-weight evidence. A fellow working this seam studies the *estimator*, not just the estimate: what the capture team actually attends to, how it updates on new intelligence, and where its confidence is miscalibrated.
- **Game theory.** The federal competition is a strategic game with a published rulebook and an observant rival. Game theory gives the discipline its sharpest vocabulary for teaming, pricing posture, and the winner's curse. A fellow working this seam runs the auction-theory results against real award data — does the theory's prediction about bid shading show up in

realized margins?

- **Behavioral economics.** Ghost themes, anchoring in price-to-win, the sunk-cost pull of a losing pursuit — these are behavioral phenomena with a research literature behind them. The fellowship’s job is to connect that literature to the doctrine’s claims and, where they conflict, to let the evidence speak.
- **Organizational learning.** The cosmology loop promises that a Won-World re-seeds Discovery. Organizational-learning research asks whether organizations actually learn from outcomes — and the outcome-data record finally makes that question testable at scale.

The cross-disciplinary mandate is a **two-way street**: the neighboring field gets a rich, public, real-world dataset, and the doctrine gets sharper questions and methods.

7. The fellowship’s expectations — a standing contract

The fellowship has no seat-time requirement and no final exam. It has a **standing contract** that is reviewed annually (see [refresh-and-recertification.md](#)):

Expectation	The bar
Research	At least one study designed and one analysis executed per year; at least one article submitted to a peer-reviewed venue per two-year cycle.
Corpus	A standing curation assignment and, over time, authored doctrine grounded in tested claims.
Mentorship	Active service on at least one doctoral committee, plus supervision of a research apprentice.
Teaching	Regular instruction at the doctoral or forever-learning altitude.
Integrity	The non-fabrication and replication benchmarks, applied to everything — the corpus, the seminar, the advisory role.

The annual review is a **gate with written criteria** ([doctrine/04](#)), the discipline’s own governance applied to the researcher. It exists so a fellow never discovers at year five that the program expected something different than they delivered — the same lesson the doctoral overview teaches about the dissertation chair.

Self-check

1. What is the difference between *testing* the doctrine (doctoral) and *stewarding* it (fellowship), and which responsibilities make stewardship concrete?
2. Why must a corpus author write the *concept* before the machine-side *encoding* — and what does the curator protect against?
3. How does a fellow's annual review mirror the discipline's own gate doctrine?