

QUANT FINANCE PLAYBOOK

ROLES / EVIDENCE / PREPARATION

The Quant Career Field Guide

Choose a role, present your work, and prepare for the conversation
that follows.

RESEARCH / TRADING / ENGINEERING

Make your evidence legible.

HOW TO USE THIS BOOK

The Quant Career Field Guide

Ten chapters, three fictional candidate cases with complete resume content, original technical dialogues, project and application worksheets, and a six-week preparation plan.

Work from your own evidence. Read the fictional examples, complete the corresponding worksheet, and retain the facts behind every claim in your application.

Original educational examples

The exercises and candidate scenarios are original teaching material. They are not leaked employer questions, employment outcomes, or descriptions of actual trading performance. Sources identify public guidance and further reading; cited organizations do not endorse this publication.

Edition 01

September 2026. Independent educational material. This edition has not been reviewed by an identified industry practitioner. Use the supplied companion files alongside the chapters, and distinguish fictional examples and synthetic results from your own evidence.

FROM INTEREST TO INSPECTABLE WORK

ROLE

Choose the responsibility

EVIDENCE

Name your contribution

EXPLANATION

Defend the claim

A dated hypothesis, a factual document, and work you can explain.

CONTENTS

Your route through the work

01 / Choose the work you want to explain	4
02 / Read a listing as evidence	6
03 / Inventory what you can defend	9
04 / A resume casebook	12
05 / Finish one coherent project	17
06 / Make your reasoning available	19
07 / Explain research outside your subfield	21
08 / Run a search you can manage	23
09 / Six weeks with observable outputs	25
10 / Use the next piece of evidence	28

CHAPTER 01

01 / Choose the work you want to explain

“I want to work in quant” does not yet tell you what to put on a resume or practice on Tuesday evening. The label covers people who investigate predictive relationships, make trading decisions and build research or trading systems. Those responsibilities overlap in some teams and are separate in others.

Which problem would you investigate after the first attractive result disappears? Start there. Your answer can change; an explicit first hypothesis gives preparation a direction.

Three problem families

Research inference. You are trying to decide what can be learned from incomplete, noisy or dependent observations. The work may include defining a target, constructing features, estimating a relationship, checking an evaluation design and deciding whether a result survives plausible alternatives. Your evidence should make those decisions visible. Naming a sophisticated model is a small part of that story.

Trading decisions. You are choosing an action under uncertainty and constraints. Information arrives, prices or opportunities change, and an existing position affects the next decision. Your preparation should include probability reasoning, explaining a decision rule, and adapting it when assumptions change. Do not infer from this description that every trader follows the same interview format or requires the same degree.

Engineering systems. You are building software with requirements that have to survive actual use. Depending on the team, the relevant work might concern research infrastructure, data quality, performance, reliability or trading systems. A project should show how you reasoned about correctness, interfaces and tradeoffs. A toy finance theme does not automatically improve an otherwise strong engineering example.

These are working categories for planning. Read the specific job description before treating one as the employer's definition. “Research engineer,” for example, can place very different weight on research methods and production systems across teams.

Employer evidence has boundaries

Jane Street's interviewing overview separates preparation by discipline and describes overlap between research, trading and engineering. Citadel's campus research guidance includes programming and discussion of research. These public descriptions support preparing both technical work and its explanation. They do not establish an industry-wide sequence of interview rounds. [Jane Street](#) · [Citadel](#), accessed September 23, 2026.

You can apply without choosing a permanent career identity. Explain why this role interests you and what evidence makes it a reasonable next step.

Three fictional candidates we will follow

Maya Chen is a final-year mathematics student. She has a probability course, a Python project on synthetic time series and experience tutoring. She has not worked at a trading firm and has no real trading results. Her preparation problem is turning academic and project work into concrete evidence without inflating it.

Daniel Okafor is finishing a physics PhD. He has developed estimation code for a noisy experimental system, collaborated on a paper and presented negative findings that changed the next experiment. He can explain his narrow research area in detail but tends to bury his own contribution inside a long literature introduction.

Leah Patel is a software engineer with three years of experience on data-processing services. She has investigated reliability failures, improved a pipeline and worked with internal users. She is interested in engineering roles near quantitative research. Her strongest evidence is already technical; she needs to connect it to the target role without inventing financial impact.

All three are invented teaching cases. Their names, institutions, projects and numbers are not customer stories or verified hiring outcomes.

A first role hypothesis

Maya chooses trading and junior research listings as initial areas to investigate. Daniel begins with research roles that value independent investigation. Leah begins with research-platform and software roles. None of these choices follows from the degree or title alone. The next chapter will compare what actual listings ask for with what each person can demonstrate.

Write your own hypothesis in two sentences: “I am initially targeting [specific responsibility] because I can already show [evidence]. My largest uncertainty is [gap], which I will investigate through [bounded action].”

For Leah: “I am initially targeting research-platform engineering because I can show reliable data-processing work and explain performance tradeoffs. My largest uncertainty is the team's expected statistical depth, which I will investigate through the listing and recruiter conversation before choosing additional preparation.”

That uncertainty directs Leah's preparation toward the role's actual requirements.

Exercise: distinguish interest from evidence

Consider these statements: “I follow financial news”; “I built a parser and documented its failure behavior”; “I enjoy hard problems”; “I compared a forecast with a baseline on later observations.” Which can a reviewer inspect?

Discussion. The parser and forecasting study point to artifacts and technical questions. Connect an interest to work: what question did the news raise, what analysis did you do, and what did you learn? Keep genuine interests; distinguish them from evidence of technical contribution.

Your output from this chapter is one initial role hypothesis and one uncertainty to resolve. Keep both dated. Later chapters should improve that choice, not merely decorate it.

CHAPTER 02

02 / Read a listing as evidence

A job description mixes several kinds of information: work the team needs done, capabilities used to do it, eligibility requirements and attractive preferences. Treating every phrase as an equally rigid condition produces two opposite mistakes. One candidate rules themselves out unnecessarily. Another ignores a plainly stated requirement because a different employer was flexible.

Read one current listing slowly. Save its URL, title, office and access date. Record the wording of any requirement that materially affects eligibility in your private notes, and distinguish it from your interpretation. A screenshot or saved copy can help you remember what was advertised, but do not redistribute an employer's full page as your own content.

Four columns that change the application

Column	What belongs there	What you do with it
Work	Responsibilities and problems	Select evidence that resembles the work
Capability	Methods, tools and communication	Identify examples you can explain
Eligibility	Explicit degree, location or other stated conditions	Confirm your fit or ask a precise question
Preference	Helpful background or desirable experience	Weigh gaps without inventing certainty

If the listing is ambiguous, preserve the ambiguity. A recruiter can clarify a relevant question; a confidence score you made up cannot resolve it.

A real contrast

On September 23, 2026, Citadel's US PhD-graduate quantitative researcher listing explicitly asked for a relevant PhD and independent research/programming evidence. Jane Street's NYC software engineer listing emphasized programming and collaboration while stating that prior finance or experience in particular named languages was not required. These are different jobs with different stated requirements. [Citadel listing](#) · [Jane Street listing](#).

The sensible conclusion is to inspect each role. It would be unjustified to infer either that all quantitative jobs require a PhD or that degree requirements never matter. A broad career slogan loses the distinction you need for a real application.

Completed fictional role matrix

The rows below are invented role briefs used to demonstrate the method. They are not current openings.

Role brief	Relevant evidence	Gap / question	Next action
Graduate research, time-series evaluation	Maya's synthetic project and statistics coursework	Little independent research beyond course scope	Prepare one self-directed extension and explain its limits
Research scientist, experimental inference	Daniel's estimation method and controlled comparisons	Unclear production-programming expectations	Ask about the balance of investigation and implementation
Research-platform engineer, batch data services	Leah's pipeline ownership and incident analysis	Unclear distributed-system scale	Compare stated responsibilities with one detailed architecture case

Notice what the matrix does not contain: an invented percentage match. Evidence differs in quality, and one explicit eligibility requirement can matter more than several keyword overlaps. A numeric total often hides those differences.

Translate requirements into questions

If a listing mentions time-series analysis, ask yourself what you can explain about time ordering, dependence and evaluation. If it mentions performance-sensitive C++, ask what you can demonstrate about the relevant language and system behavior. If it asks for collaboration, identify an actual disagreement, review or handoff and your contribution to it.

Do not add a tool to the skills section because its name appears in the listing. That creates a question you may be unable to answer. If you have adjacent experience, describe the adjacency directly: “Production Python data services; currently learning C++” is different from claiming equivalent production expertise in both.

A bounded clarification

For a role whose title is unclear, a concise question might be: “The description includes both model investigation and research infrastructure. How is the team's work divided between those responsibilities for this position?” It asks about a decision that affects fit and preparation. It does not ask the recruiter to design your whole career plan.

Requirements about location, start date or work authorization belong with the specific employer. Record your own facts accurately and ask the appropriate contact when the listing is unclear. This guide does not infer eligibility from your degree or passport.

Exercise: classify five statements

An invented listing says: build reproducible pipelines; strong Python; relevant doctorate required; familiarity with C++ useful; collaborate with researchers. Sort them into the four columns.

One reasonable classification. Pipeline construction and collaboration describe work. Python is a capability. The doctorate is an explicit eligibility requirement. C++ is a preference in this wording. Collaboration also implies a capability worth demonstrating. Categories can overlap; the purpose is to preserve the statement's force, not win a labeling contest.

Keep a short decision record

For each priority role, retain one sentence explaining why you will apply, defer or investigate further. Revisit it when new information arrives. A large spreadsheet with no decisions becomes clerical work; a small matrix that changes your next action earns its place.

CHAPTER 03

03 / Inventory what you can defend

Before writing a resume, collect the facts from which it can be written. Memory compresses a project into its most flattering conclusion. An evidence ledger forces you to separate the question, your action, the result and the limits of that result.

For each experience, record what existed before your involvement, what you personally changed, how you evaluated it, and what artifact supports the claim. Add a boundary: what would be overstated? This last field is especially useful in group research, where the team's accomplishment and your contribution are related but not identical.

Maya's project ledger - fictional case

Maya completed a course project using synthetic time-series data. Her group had three members. She wrote the evaluation split logic, compared a linear forecast with a constant baseline and found an accidental future-valued feature. Another member wrote the initial generator and another prepared the presentation.

Field	Recorded fact
Question	Does a current synthetic state help predict a later outcome?
Personal work	Time-ordered evaluation, baseline comparison, feature-availability review
Team work	Generator, report and presentation shared among three students
Evidence	Commit history, notebook, short written explanation of the removed feature
Result	Corrected the invalid comparison and documented the revised result
Boundary	No live-market data, trading profit or sole ownership of the whole project

She can write a useful bullet without claiming a return. Finding and repairing an invalid comparison is substantive work. If a final numerical metric is missing from her notes, she should retrieve it or omit it, not replace it with a plausible value.

Daniel's research ledger - fictional case

Daniel's research group studied a noisy physical measurement. He implemented an estimator and compared it with a simpler method across 12 controlled simulation settings. In four low-signal settings the more elaborate estimator was unstable. The group changed the proposed experiment after reviewing those failures.

His ledger distinguishes implementation from theory developed jointly with an adviser. It records what “12 settings” means, how randomness was controlled, which metric was compared, and whether the simulation

code can be shared. A paper citation establishes a publication, but does not on its own explain his personal technical contribution.

A defensible account might be: “Implemented and evaluated an estimator across 12 controlled simulation settings; identified instability in four low-signal settings and documented conditions favoring the simpler baseline.” Those numbers are fictional inputs for this case. They are never a suggestion to add similar numbers to a real resume without evidence.

Leah's engineering ledger - fictional case

Leah helped maintain a daily data pipeline. She redesigned a retry mechanism after duplicate processing caused inconsistent downstream outputs. Her contribution included an idempotency key, failure-state logging and a migration plan reviewed by two teammates. The team observed a reduction in duplicate records over a defined monitoring window.

The useful facts are the failure mode, her implementation, how behavior was checked and the scope of observation. If the measured reduction cannot be reliably attributed to her change, the bullet can describe the technical improvement and observed behavior without claiming sole causality.

She should not translate “data pipeline” into “trading infrastructure” because the target firm trades. The connection is the engineering problem: correctness under retries, observable failure states and reliable interfaces.

Four kinds of numbers

Scale describes what you handled: rows, users, experiments or services. It is useful only with context. A million synthetic rows may be easy to generate and say little about research quality.

Comparison describes a measured difference from a named baseline. Include the evaluation setting and units. “Reduced error” needs an error definition.

Reliability describes behavior over a specified period or set of cases. Avoid turning a short observation window into a claim that failures can never recur.

Scope of ownership describes your part: implemented two components, led a review, reproduced one result. It prevents a team outcome from being misread as individual authorship.

No rule requires a number in every bullet. A precise methodological or engineering contribution can be stronger than an arbitrary count.

Exercise: repair an unsupported claim

Original: “Generated significant alpha with a machine-learning model.” Available evidence: a course notebook, synthetic data, a train/test comparison and no execution model.

Repair. “Evaluated a forecasting model on synthetic data using a time-ordered split and a constant baseline; documented the difference between prediction error and a tradable outcome.” Add a verified result only if the

record supports it. The correction preserves the work while removing an economic claim the evidence cannot establish.

Ownership questions to answer now

Which function or analysis did you write? What decision did you make when alternatives existed? What did a collaborator contribute? What result did you initially misunderstand? Which files can you share? If you cannot answer, revisit the project record before polishing the prose.

The completed ledger becomes the source for your resume, project explanation and interview preparation. Update the facts once and derive each presentation from them. That is easier to maintain than inventing a slightly different story for every format.