# AI in P3 preconstruction: what changes for bid teams

This is a Mirage Metrics white paper for the people who prepare public-private partnership (P3) and progressive design-build bids in Canada and the United States. A P3 bid packs a project agreement, its schedules, a reference design, several design packages and a stream of addenda into a few months, and splits the work between partners who each keep their own estimate. AI agents now take over the reading that repeats. They register requirements against their clauses, measure quantities from each package and tie them to their sheet, list the changes from one addendum or one package to the next, and check the documents against each other. People keep what a bid is made of: the design choices, the price, the risk, the decision to bid.

The paper states no figure of its own: rules and procedures are on the linked guides, with their primary sources, and Mirage's figures are those of its published case studies.

https://miragemetrics.com/research/ai-in-p3-preconstruction

## The short version

- A P3 or progressive design-build bid is a document problem before it is a pricing problem: the project agreement and its schedules, a reference design, design packages at increasing maturity, addenda, and partners' estimates that must agree.
- The work AI agents do well is the reading that repeats: registering each requirement against its clause, measuring quantities from each package and keeping each tied to its sheet and revision, listing what an addendum or a new package changed, checking documents against each other.
- What stays with people is everything a bid is judged on: design choices, alternative technical concepts, prices, risk allocation, the legal reading of the agreement, the decision to bid.
- Inside a joint venture, where the agents run and who holds the data matter as much as what they do: the deployment, the access and the audit trail are part of the evaluation.
- The safe way to test is on a bid the team has already submitted, against what the team produced by hand.

## Where a P3 bid team's weeks go

The procurement differs by owner and by model, and the guides linked here describe each in detail. The shape of the work, however, repeats. After the request for qualifications, the shortlisted teams receive a request for proposals with a draft project agreement whose schedules carry the owner's requirements: the output specifications, the payment mechanism and its deductions, the design and construction certification procedure, the variation procedure (see how bid teams review P3 project agreements). Under a progressive model, the design and the price mature together over a development phase before a price is fixed (see project delivery methods).

Five kinds of work take the time. The first is reading. The agreement and its schedules run to thousands of pages, and every requirement that carries a cost, a quantity, an interface or a deduction has to be found and assigned. The second is measuring. Quantities are taken off the reference design, then again at each design package, and every move between packages has to be explained, which is what an estimate's class depends on (see AACE estimate classes). The third is tracking, because each addendum changes documents already read and quantities already

priced. The fourth is reconciling: the partners of a joint venture each estimate their scope, and their quantities must meet at the interfaces. The fifth is pricing options, since an alternative technical concept means re-measuring a part of the design in days.

## What AI agents do on a P3 bid, task by task

The table separates what an agent produces from what a person decides on it. Each line links the guide that describes the documents involved.

| Task | What the agent produces | What a person decides | Where the documents are described |
|---|---|---|---|
| Project agreement review | A register of requirements, each with its schedule and clause, its owner on the bid team, and the other clauses it refers to | What each requirement costs, who carries its risk, how it is priced or qualified | P3 project agreements |
| Addenda | The list of what each addendum changed against the previous version: clauses, drawings, quantities | Which change moves the price or the design, and the questions to send the owner | P3 project agreements |
| Quantities | Quantities measured from the drawings and schedules, each tied to its sheet and revision, structured against the bid items | The means and methods, the production rates, the unit costs | AI for heavy civil estimating |
| Design packages | From one package to the next, the quantities that moved, with their sheets in both packages | Which moves are scope growth, design development or an error, and how the estimate's class changes | AACE estimate classes |
| Owner specifications | The measurement and payment rules of the owner's specifications applied to each quantity | How a measurement rule is read where the drawings and the specification disagree | OPSS (Ontario) |
| US federal transit cost structure | Quantities mapped to the FTA Standard Cost Categories and kept in sync with each package | The category of an item that belongs to two, the contingency | FTA Standard Cost Categories |
| Historical prices (US) | Bid item quantities set against the owner's published bid tabulations | Which historical price applies to this job, in this market | DOT bid tabs |
| Bridges across the border | Quantities read from drawings in metric or imperial units, with the code each structure is designed to | What a different code changes in the design, the cover, the deck | CSA S6 vs AASHTO LRFD |
| Québec bids | The bordereau checked against the drawings and the ministry's specifications, in French | The price of each item, the conformity of the soumission | Soumission au MTQ (in French) |

Two properties make an agent's output usable on a bid. First, every line is traceable to the document it came from, so a reviewer can check it in seconds rather than redo it. Second, the output follows the structure the team already prices in (bid items, pay items, cost codes), so it lands in the estimate instead of beside it. The systems those estimates live in, and how each one takes quantities in, are on the estimating software integrations page.

## What stays with people

People keep everything a bid is judged on. That covers the design solutions and the alternative technical concepts a team proposes, the price of each item and the margin, and the allocation of risk between partners, designers, suppliers and the owner, with the qualifications that go with it. It also covers the legal reading of the project agreement and the decision to bid at all. An agent that measures and reads faster changes how much of the bid period those people spend on these questions, but it does not answer them.

Nor does an agent replace the engineers' work on the design itself: structural design, the erection sequence of a bridge or a guideway, the tunnelling method. It measures and reads what that work produces.

## Data and control inside a joint venture

A P3 bid team is usually a consortium whose partners compete with each other on the next pursuit. Three questions decide whether an AI system can be used at all. The first is where it runs. Mirage can run in the team's own infrastructure or cloud account, in which case the data and processing stay in the environment the team controls and the encryption keys are its own. The second is who sees what. Access is set by role, so each partner sees what the consortium agreement allows. The third is what is recorded. Every agent action is logged with its inputs, its sources and its result, and nothing is written into a partner's systems before a named person approves it. The commitments, the deployment options and the answers to IT teams are gathered in the procurement and security pack, which a bid manager can forward to each partner's reviewers.

## How to test it without risking a pursuit

- Pick a bid the team has submitted: its agreement, its packages, its addenda and the quantities it produced by hand.
- Run the proof of concept on it, one week after the first meeting, and set the agent's register, quantities and change lists against the team's own.
- Judge it on traceability and on what a reviewer has to redo, not on a single accuracy figure.
- If it holds, put one workflow into production on a live pursuit: the first workflow is in production in 4 to 6 weeks.

## Where Mirage Metrics fits

Mirage Metrics builds and operates AI agents for construction estimating and preconstruction, deployed by its own engineers inside the client's systems. The homepage names Flatiron Dragados, Demathieu Bard Construction, Green Infrastructure Partners and Skyway Bridging Partners among the companies it is in production at. This paper says nothing about their work with Mirage. The published construction results are OCBM's, where three agents run in production with every quantity tied to a source drawing, and John Cockerill's, where piping takeoff takes 80 to 90% less time (OCBM, John Cockerill). The North American projects these teams bid are on the infrastructure map.

## Questions readers ask