WORKING PAPER NO. 1

We Commit Before We Know

A working paper on the Production Commitment Point in manufacture-and-install businesses

Brian Milburn

August 2026

CONTENTS

Why I'm writing this

I've spent my whole career in construction, on the manufacturing and installation side of it. I keep hitting the same problem from different directions. I've never found it written down in a way that matches what I see. So I'm writing it down myself, clearly enough that people who know more than I do can tell me whether it holds up.

Throughout this paper I use the term **manufacture-and-install business** to describe companies that design, engineer, manufacture, and install custom-built products into buildings that are still under construction. This paper is about that operating model, not about any one trade.

Let me be clear about what this is not. It isn't a knock on existing operations theory. Lean, Quick Response Manufacturing, the Last Planner System, and the Theory of Constraints are each mostly right for the job they were built for. I use all four.

The question I'm trying to answer is a different one. What's the best way to describe a company that both manufactures a custom product and installs it into a building that's still being built?

I don't think that question has a settled answer. I could be wrong. If somebody has already answered it, point me at their work and I'll stop here.

Here's the whole thing in one paragraph. A manufacture-and-install business runs two production systems at once. We build in a shop we control. We install into a building other people are still building. We watch our own shop closely and all the time. We usually check that building loosely and once in a while. That difference creates a blind spot in how we decide when to commit work into production. The cost of it shows up later, somewhere else, on somebody else's line.

Everything below is an attempt to say that carefully enough to test.

---

1. Defining the system

Before I say what might be missing, I should say what I think we're actually running.

These businesses run a mixed system with at least eight phases. Estimating. Drafting and design. Engineering. Procurement. Production. Delivery. Installation. Closeout. Long-lead procurement usually runs parallel to the front end rather than waiting its turn.

None of that is new. Every shop in this business runs the same list.

What's worth noticing is who owns each phase. Estimating has an estimator. Production has a production manager. Installation has a field lead. Each one owns a box and gets measured on that box.

The project manager is the only role that carries the job through all eight. Not running them. Carrying the thread through them.

That means the project manager is the only continuous part of a system made of parts that hand off. I'll come back to that.

Here's what I think makes it different from ordinary manufacturing:

The place that receives our product is still being built.

A normal manufacturer ships to a warehouse, a dock, or a store. That place sits still. It was the same last month and it'll be the same next month.

We ship to a building under construction. Other people keep changing it, on their own schedules, for their own reasons, the whole time we're building.

So we have two production systems running side by side. Ours, and the site's. There's one handoff between them, and we only control one side.

The two commitment points

Manufacturing already has a name for the moment work stops being generic and becomes this customer's job. In engineer-to-order work, which is what we do, that moment sits way back at design. That part is settled and I'm not claiming anything new about it.

What I want to suggest is that we have a second commitment point later on. That's the moment the product has to fit a real opening in a real building. I'll come back to the first one, the point where our own money goes in, and give it a name.

Weeks pass between those two points. The opening isn't holding still.

I want to be careful here. Several bodies of work already cover pieces of this. I'm building on them, not filling in blank space.

- **Decoupling point research** looks at where in the process a job stops being generic and becomes this customer's job. Some of that work says you should hold off on detail drawings until the site and the shop actually need them, instead of drawing everything up front. For example, instead of fully engineering every room six months before production, you only complete the rooms the shop actually needs to build next.
- **Design freeze practice** in offsite and modular construction deals with what a change costs once the shop has already started building. Some of that work looks at locking the design in stages instead of all at once, so the parts the shop needs early get locked early and the rest stays open longer. For example, on a hospital project you may release long-lead laminate, sheet goods, and hardware for procurement while leaving final filler widths or scribe details open until field verification. You're freezing only the parts of the design that must move today.
- **Tolerance research** looks at how small errors add up. The shop is off an eighth, the wall is off a quarter, the floor is out of level, and together they turn into a piece that won't fit. Some of that work can predict the odds of a fit problem before you build. It's the idea that nobody made one big mistake, but several small ones stacked together into a major fit problem.
- **Make-ready planning** in the Last Planner System says don't schedule work until you've gone looking for whatever would stop it and cleared those things out of the way. There's also work out there on running a factory to what the site can actually take, instead of running it flat out. For example, if drywall isn't finished or another trade still owns the room, make-ready planning says don't send your installers there just because the schedule says it's time.
- **Value of information** research asks what it's worth to go find something out, compared to what it costs to find it out. Sometimes the answer is that you're better off deciding without checking. That's a real answer, not a failure. For example, if a one-day site visit catches missing blocking before drywall goes up and prevents remaking $40,000 worth of millwork, the value of that information wasn't the cost of the trip. It was the cost of the mistake that never happened.

What I haven't found is anybody asking what happens when the chance to check expires. Here's what I'm suggesting, and it's narrow:

> **Proposal:** Some site conditions can only be checked during a window, or can't be answered by a single check at all. We commit work against those conditions anyway, because checking costs money now and the payoff lands somewhere else, later.

A room measured six weeks ago usually hasn't changed. Most walls stay put. The problem isn't that the measurement went bad. The problem is that nobody went back to see which rooms did change.

Blocking is the clearest case. You can see blocking after the framer sets it and before the drywall goes up. That's the window. After the rock is on, you can't check it without opening the wall.

Miss that window and you don't have bad information. You have no information, and no way to get it.

Nobody misses it out of carelessness. Checking means sending somebody out during a specific two week stretch to look at a wall, on a job that doesn't install for two months. That trip costs a person, a day, and fuel. So it doesn't happen. It doesn't happen for good reasons, every time.

Two kinds of conditions

Not everything fails the same way. There are at least two kinds.

Some conditions have a window. Blocking, backing, substrate, rough-in behind the finish. One look at the right moment answers the question for good. Miss the moment and there's no answer available at any price.

Some conditions are a history, not a state. Conditioned space is the example everybody knows. Somebody walks the room Tuesday and it's 72 degrees at 45 percent humidity. That tells you nothing about the three weeks before or the two weeks after. And by the time you can see the result in the product, the panels have already cupped and the doors have already moved.

Those two are different problems, and the difference shows up in cost.

A condition that's a history can be measured by a device. It costs real money to put a sensor in a room. After that, checking it a thousand times costs about the same as checking it once.

A condition that has a window can't be handled that way. There's nothing to set up. Every check costs another trip, on every room, on every job, and that cost never comes down.

There's a second reason those don't get solved the same way. Checking blocking means checking an absence against an intention. You aren't looking at a thing. You're looking for whether the thing that should be there is there, at the height the drawing says, in the bay the drawing says, heavy enough for what's going on it. That takes somebody who knows what was supposed to happen in that wall. The missing piece doesn't photograph.

I'm not recommending anything here. I'm saying the two have opposite cost structures, and that may be why one keeps getting attention and the other doesn't.

What the code already knows

Building codes already work this way. Framing and rough-in inspections exist because covering the work makes it impossible to check afterward. The code sends somebody to look before the wall gets closed.

We accept that. Nobody argues about it.

Then we commit a hundred thousand dollars of product against that same wall without ever having gone to look at it ourselves.

The code recognizes that some things have to be checked before they're covered up. Our own operations don't.

One more thing worth saying

Conditioned space is usually in the contract. The GC owes us conditioned space before delivery, and the GC tells us the space is conditioned.

Nobody can check whether the heat stayed on over a long weekend three weeks ago.

That's a binding condition, reported to us, that can't be audited after the fact. When it's wrong, we find out from the product.

Defining the Production Commitment Point

I used to call this the release point. I've stopped, because "release" already means four other things. Engineering release. ERP release. Production release. Schedule release. Those are all administrative acts, and none of them is what I'm talking about.

What I'm talking about is the moment we're economically on the hook.

> **Production Commitment Point (PCP):** The point at which labor, material, and production capacity become committed to a specific room or product such that reversing the decision becomes increasingly costly.

Where that point sits depends on the shop. It might be first cut. It might be material hitting the saw, the nest, or the brake. It might be the moment a batch goes into a run, or the point where you order material against a job.

Which step it is doesn't matter much. Before the PCP, changing your mind is cheap. After it, changing your mind costs you.

Every shop has that point. Most of us never name it. That's part of the problem.

From here on I'll use Production Commitment Point, or just commitment point where it reads better.

What else feeds the commitment point

Site conditions are one input out of many, and I'm not suggesting they replace the rest.

A commitment point fairly weighs shop capacity and current load, backlog, customer commitments and promised dates, procurement status, material lead times, crew availability, cash needs, and batching efficiency. Those are all real. Most of them already get managed well.

A shop that committed work purely on site readiness and ignored everything else would be badly run.

My argument is narrower. Out of everything feeding that decision, what's happening at the site is usually the least checked. It's also the one thing that decides whether the work will fit when it gets there.

The system, drawn

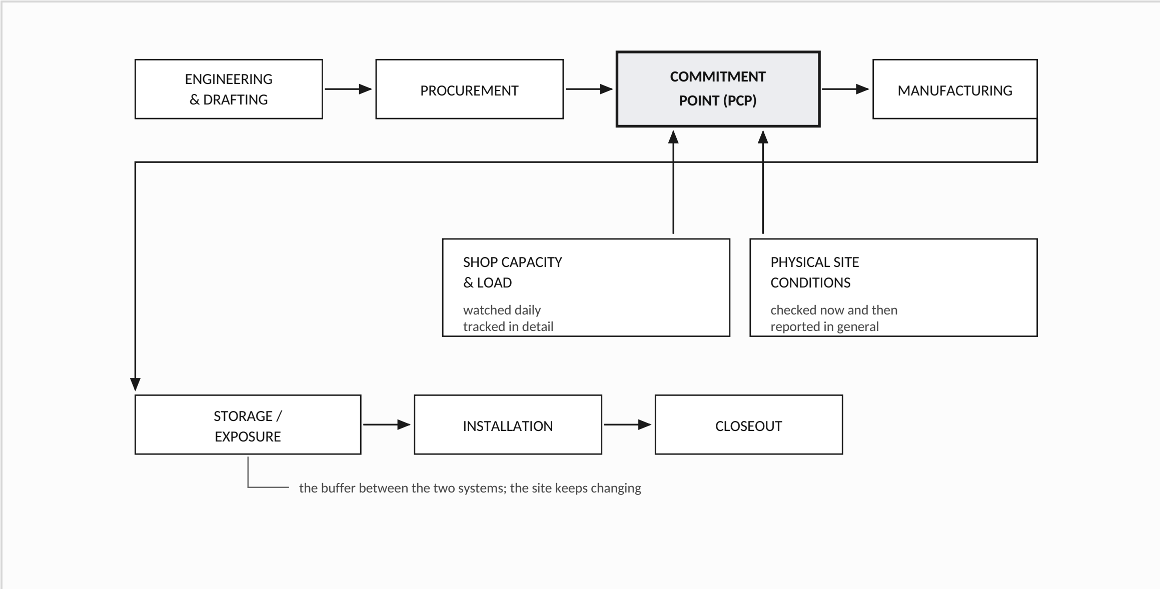

One decision, two inputs, and we know far less about one of them.

---

2. Why we commit early

I want to head off a reading I don't intend. Nothing here says shops should just hold the commitment point longer. That isn't how this business runs. A paper that said so would deserve to be ignored.

We're cash-intensive. We stay alive by keeping engineering, procurement, manufacturing, and production moving.

There are good reasons to build and store ahead of a site being ready. It brings in revenue. It keeps production flowing. It frees up capacity for the next job. It meets commitments we already signed.

A lot of the time, committing early isn't a compromise. It's the only practical call on the table.

We also handle part of this risk in the contract. Approved hold-to dimensions push ownership of some size changes onto somebody else. That's real and it matters. It changes who pays.

But it doesn’t stop the disruption if things move later. The crew still stands around. The piece still comes back to the shop. The next job still gets pushed. We may get paid for the change and still lose the week.

There may be something bigger underneath all of this. I’ll put it out here without pretending I’ve worked it through.

Maybe this problem sticks around for a reason other than bad calls by project managers. Maybe the money side of this business more or less forces us to split building from installing.

A plumber or a framer builds in place. We build a custom product weeks or months before anybody installs it. So inventory and storage end up acting as the cushion between our shop and the jobsite.

If that’s right, the gap between building and installing isn’t a scheduling failure. It’s built into how we get paid and how our capacity turns over.

The same pressure works on checking. Every trip to a site we aren’t installing on yet is a cost we carry now against a benefit that lands later, on a different line, in a different month.

I don’t know whether that belongs in this paper or a different one. My guess is it deserves its own. But it changes the question I’m asking, so it needs to be on the table.

Here’s what I keep coming back to. The question isn’t whether we should commit early. Often we have to. The question is whether we know what we’re taking on when we do.

---

3. What “readiness” actually means

I’ve used the word “readiness” loosely for years. It needs to be broken apart. Lumping it together is part of why it’s so hard to manage.

There are at least five different kinds:

- **Design readiness.** Drawings approved, RFIs answered, owner decisions made.
- **Procurement readiness.** Materials and hardware confirmed and available.
- **Manufacturing readiness.** Capacity, sequence, and shop capability available.

- **Site readiness.** The room itself is ready. Wall surface, blocking, dimensions, conditioned space, trades ahead of us finished.
- **Installation readiness.** Crew, access, logistics, staging, permission to be there.

Field readiness is one piece of project readiness, not another word for it. That matters for a claim I make later, so I'm saying it up front.

---

4. The observation

Here's what I think I can state as something I've seen, not something I'm arguing. It comes from my own experience. I'd like it tested against a wider sample before anybody calls it an industry fact.

> **Observation:** In the shops I've worked in, visited, or dealt with, we run the shop on numbers that get updated constantly. We run site readiness on conversations. Phone calls, emails, meetings, a visit now and then, and whatever the project manager has learned over the years.

Ask a shop of any size for machine hours by department, load by week, backlog, bottleneck. Most can hand it to you.

Then ask the same shop how it knows a room is really ready to receive finished product. In my experience the answer is a call with the superintendent, plus a project manager's read on whether this GC's schedule means anything.

I don't want that to sound dismissive. Experienced project managers are very good at that read. Their judgment is a real skill, not a workaround. But it lives in people instead of in the system. Nobody measures it, so it doesn't get better on purpose.

Here's the plainest way I know to say it. The shop knows what the machines are doing because the system tracks it. The field often works because experienced people are constantly connecting the dots through phone calls, texts, meetings, memory, and judgment. In effect, the project manager becomes part of the operating system.

This is where the eight phases come back. Everyone else hands off. The project manager never does. So every gap between two phases lands on the one role that touches both, because there's nowhere else for it to go.

The advantage in this business has always lived in the connections between the phases rather than in the phases themselves. That idea isn't mine. Michael Porter made the case for it in 1996, and Marc Sanderson made the same point about our industry recently. What I'd add is narrower. In a manufacture-and-install business, some of those connections aren't systems at all. They're a person.

Both of those inputs feed the same decision.

One side gets watched all the time and in detail. The other gets checked once in a while and comes back rough. It isn't a difference in honesty or effort. Nobody is hiding anything. We're just watching one half of our own operation a lot closer than the other half.

---

5. The inference

From that I'd draw a conclusion. I'm labeling it a conclusion because I can't prove it.

> **Inference:** We reach the Production Commitment Point with good shop numbers and rough site information. So the decision may not account for what's really happening at the site. What that costs us shows up later as rework, delay, storage, and field changes. It never gets charged back to the commitment that caused it.

Here's what I find interesting. The cost of checking lands now, on us, on a job that isn't installing yet. The cost of not checking lands later, on the crew, on a different line, in a different month.

So the trip never gets made, and nobody is wrong for not making it.

There's a second conclusion about order, and this is the one I keep circling back to.

> **Inference:** In this system, we commit before we know. The Production Commitment Point often arrives when nobody has gone and looked at that room in weeks. For some conditions, the chance to look has already passed.

If that's right, it's backwards from how you'd want to order any decision this expensive.

I'm not saying anybody designed it that way. It looks like a side effect of two systems that grew up apart. Manufacturing planning got very good. Construction planning got a lot better. The handoff between them never got worked out on its own, because it doesn't sit fully inside either one.

---

6. The hypothesis

> **Hypothesis:** Rooms committed after somebody checked the conditions that matter, while checking them was still possible, will come out better than rooms committed without that. Better meaning fewer rework hours, fewer field change hours, and more dates hit.

Three limits on that, and I put them there on purpose.

First, this is only about conditions we can't answer later. Blocking, backing, substrate, rough-in behind the finish, and conditions like temperature and humidity that are a history rather than a state. Room dimensions don't work the same way, since you can always go measure a wall again. The claim is strongest where covering the work, or letting time pass, ends our ability to know.

Second, this is about odds, not cause. Better site information doesn't cause better outcomes. It improves the odds by taking one known source of error out of a decision that has plenty.

Third, this is about the room, not the job. Jobs don't get ready. Rooms do. Most scheduling systems still treat a job with one room ready and one room six weeks out as a single unit.

The room is probably the better unit for commitment. I hold that one more loosely than the rest. It may not hold up on jobs with big shared runs, where breaking up the batch costs more than the risk it saves you.

7. On measurement

If this idea is worth anything, it ought to be measurable.

The number I'd start with is how many rooms we committed without anybody having checked the conditions we can't go back and check. Call them unverified commitments.

Any shop could count that today with no new tooling. Pull twenty rooms. For each one, ask whether anybody confirmed the blocking and backing before the wall got closed. Then ask what that room cost in field change hours.

Two more worth measuring alongside it.

What a check actually costs. Not a guess. A person, a day, a truck, against a job that isn't installing yet. That number is the reason the trip doesn't happen, and I've never seen anybody write it down.

How many windows are open right now across the active jobs, and how many are about to close. Nobody tracks that, because nobody frames the problem this way.

Shelf time, meaning the days between build-complete and install-start, is still worth watching. It's simple and most shops can pull it from data they already keep. But it's a weaker indicator than I first thought. Most rooms that sit a long time turn out fine. Shelf time tells you about exposure in general. It doesn't tell you whether anybody checked.

I don't know which of these is best. I'd like to see somebody test more than one.

8. How this could be shown to be wrong

I'd rather put this out in a form that can fail.

The observation fails if a broad survey of shops finds site readiness already gets tracked with hard data, updated as often as production data. If so, the gap I'm describing belongs to my corner of the industry, or just to me, and it doesn't hold up anywhere else.

The inference fails if commitment points already account for the site well enough through phone calls and experience. In other words, if project managers are accurate enough that going and looking would add nothing. That's a real possibility. I'd want it tested, not waved off.

The hypothesis fails if rooms where somebody checked the concealed conditions come out no better on rework or on hitting dates, once you account for everything else. Everything else means design changes, owner changes, how well we measured, drawing errors, and how well the shop built it. Any one of those could fake a connection or hide one.

A shop with good field measuring and a steady owner might carry a lot of risk with no penalty at all. That would be worth knowing on its own. It would point at how well we measure, not at when we measure.

There's also a fair argument that the checking isn't worth what it costs. If the trips cost more than the rework they prevent, then not going is the right call and the current practice is already correct. That result would settle the question, and I'd want it measured rather than assumed either way.

I should also flag the most likely trap. Shops willing to start checking the field closely are probably better-run shops to begin with. Any study has to account for that, or it'll just find what it went looking for.

---

9. Possible organizational implications

If the observation in this paper is right, the effects may reach past production numbers.

When readiness gets managed mostly through individual judgment and phone calls, experienced project managers and field leaders become part of the operating system, the way I said earlier. They aren't just running projects. They're carrying information the company itself doesn't keep.

Put plainly, our only continuous element is a person. People leave.

That may create risk in training, succession, onboarding, and passing knowledge along. It may also add to the load on jobs that are already hard to fill and hard to keep filled.

It may also feed itself. The more the system leans on individual memory, the more disruption it creates. The more disruption it creates, the more it leans on the people who are good at soaking it up.

I'm not claiming that poor readiness information causes burnout or turnover. Those have many causes. Workload, culture, pay, customer behavior, staffing levels, leadership.

What I'm suggesting is that a system this dependent on individual memory and constant firefighting may make hard jobs even harder to stay in.

That question is outside the scope of this paper. I think it's worth studying.

10. What I'm unsure of

- Whether the second commitment point is already modeled somewhere and I've just missed it. That's the one I'd most like corrected.
- Whether "conditions we can't answer later" is the right boundary, or whether it's too narrow. Blocking and conditioned space fit cleanly. Other trades finishing their work doesn't, since that stays visible right up until install.
- Whether the room is a practical unit for commitment at scale.
- Whether this is better handled by checking more or by shortening the gap between the commitment point and install. Those are two different fixes. The second one runs straight into the cash and capacity reasons we commit early in the first place. I don't know which is more workable.
- Whether the argument in Section 2 holds up, meaning the idea that our economics more or less force building and installing apart. I put it in because it changes the question, not because I've worked it out. It probably deserves its own paper.

A note on sources

This is a working paper written by somebody who runs operations, not a literature review. I want to be honest about that.

The work I've referred to is real and well established. I've read some of it all the way through and some only in part. Anything I've said about what the research does or doesn't cover should be treated as a guess until somebody in that field checks it.

Here's what I've drawn on, and what I still owe:

- Gosling, J. and Naim, M. (2009). Engineer-to-order supply chain management: a literature review and research agenda. *International Journal of Production Economics.*
- Wikner, J. and Rudberg, M. (2005). Integrating production and engineering perspectives on the customer order decoupling point. *International Journal of Operations & Production Management,* 25(7/8).
- Vrijhoef, R. and Koskela, L. (2000). The four roles of supply chain management in construction. *European Journal of Purchasing & Supply Management.*
- Ballard, G. (2000). *The Last Planner System of Production Control.* PhD thesis, University of Birmingham.
- Ballard, G. and Howell, G. (1998). Shielding production: an essential step in production control. *Journal of Construction Engineering and Management.*
- Hopp, W. and Spearman, M. *Factory Physics.* Particularly the variability buffering law, which says variability in a production system gets absorbed by some mix of inventory, capacity, and time.
- Viana, D., Formoso, C.T., et al. (2018). Planning and controlling design in engineered-to-order prefabricated building systems. *Engineering, Construction and Architectural Management,* 25(2).
- Bortolini, R., Formoso, C.T. and Viana, D. (2019). Site logistics planning and control for engineer-to-order prefabricated building systems using BIM 4D modeling. *Automation in Construction,* 98.
- Monte Carlo simulation for tolerance analysis in prefabrication and offsite construction (2019). *Automation in Construction.*

- Porter, M. (1996). What is strategy? *Harvard Business Review.* The argument that competitive advantage comes from how activities fit together rather than from the activities themselves.
- Value of information research in decision analysis. I have not confirmed whether anyone has applied it to construction verification, and that needs checking.

I left out trade magazines, vendor material, and industry websites on purpose. They're useful for seeing how people work today. They aren't evidence, and I don't want this paper resting on them.

What I'd ask of anyone reading this

If you run a shop and you've measured any part of this, I'd like to hear about it. Especially if what you measured says I'm wrong.

If you work in operations research or construction management and know of work that already covers this, point me at it. I'd rather read it than keep reasoning from my own experience.

I'm reasonably confident about the observation. I think the inference is sound but unproven. The hypothesis is a proposal.

About the Author

Brian Milburn comes from a multigenerational woodworking background and has spent nearly three decades in the millwork industry, including more than two decades in operations leadership. This paper reflects observations from that experience and is offered as a working proposal for discussion rather than as settled theory.