

**Ep #104: Quality Standards for Law Firms:
How to Make Expectations Explicit and Work Predictable
[Agile Lawyering Part 4]**



Full Episode Transcript

With Your Host

John E. Grant

[The Agile Attorney](#) with John E. Grant

Ep #104: Quality Standards for Law Firms: How to Make Expectations Explicit and Work Predictable [Agile Lawyering Part 4]

For a lot of legal professionals, one of the most persistent sources of stress and anxiety in their work is just uncertainty. Not knowing who's doing what, not knowing when something is actually done or whether it's been done correctly, not knowing whether work you've delegated is going to boomerang back at you later, right when you can least afford to deal with it.

A lot of that uncertainty gets misinterpreted as a communication breakdown or maybe a people problem. But more often than not, it's actually a system design problem. Specifically, it reflects a system in need of shared explicit quality standards. Now, in most practices, quality expectations exist, they just live in people's heads. And when quality lives in your head, then delivering quality work becomes a guessing game. Signals get distorted, frustrations mount, and rework quietly eats up the capacity you need for other things.

So in today's episode, I'm going to give you the tools you need to create a set of simple, easy-to-implement quality standards that will serve as a stabilizing force for your law practice. Once you start to adopt them, you'll quickly notice how flow becomes calmer, work gets more predictable, and your and your team's commitments will start to become more credible.

You're listening to *The Agile Attorney Podcast*, powered by GreenLine. I'm John Grant, and it is my mission to help legal professionals of all kinds build practices that are profitable, sustainable, and scalable for themselves and the communities they serve. Ready to become a more Agile Attorney? Let's go.

A quick note. The concepts I'm discussing today should be useful to you no matter what kind of law practice you're part of or what tools you use. If you'd like, stay tuned at the very end where I will briefly discuss how my software tool GreenLine helps support and reinforce the Agile practices from today's discussion.

Ep #104: Quality Standards for Law Firms: How to Make Expectations Explicit and Work Predictable [Agile Lawyering Part 4]

Welcome back, everyone. In the last couple of episodes, I've been framing a law practice as what I've been calling a promise-keeping machine. And at a systems level, that's really what we're building, a way to make promises to clients, to courts, to counterparties, and to each other, and then keep those promises reliably over time.

In episode 102, we talked about making work visible, and then last week we talked about bottlenecks, when work gets stuck and delivery debt accumulates.

Today, I want to tackle the other two major flow problems: duplicate effort and rework. Both of these are forms of turbulence and backflow. Work moves forward, but then it either gets duplicated unnecessarily, or it swirls back upstream for correction or clarification or completion. And like bottlenecks, these problems consume your and your team's finite capacity, they disrupt your planning, and they make it harder to keep your commitments.

And the good news is that bottlenecks, duplicate effort, and rework all have a common root cause: unclear expectations about the who, the what, and the when for any given deliverable.

So today, I'm going to focus on how explicit quality standards are the best way to address that root cause. And when I say quality, I want to make sure we know I'm not talking about high quality, although that is probably your brain's default. What I want to go for is fit-for-purpose quality, that Goldilocks zone where we're delivering things that are just right, not underbuilt, not overbuilt.

Remember that effective promise keeping depends on two guardrails being in place: consistency and predictability. Consistency without predictability is going to create anxiety. Predictability without consistency will lead to mistrust. You need both of these things. Today though, I'm going to focus

Ep #104: Quality Standards for Law Firms: How to Make Expectations Explicit and Work Predictable [Agile Lawyering Part 4]

primarily on the consistency guardrail, in part because consistency is one of the best ways to create predictability in the first place.

What I know from doing this work is that when quality expectations are clear, are shared, and are appropriate to the purpose of the work, then flow becomes more stable. Fewer things come back for rework, fewer surprises show up late, and then that stability is what makes your credible commitments possible.

I also want to emphasize again that we're not necessarily going straight for making the work move faster. We'll get there. But the more important thing for now is to help work move cleanly and smoothly with minimal stagnation and without that backflow. And that's what allows the promise-keeping machine to work well and then improve over time.

Now, before I talk about how to build those quality standards, I want to spend a few minutes illustrating how lack of consistency and lack of predictability manifest as problems in a typical law practice. And I actually started that in the last episode when I introduced this idea of boomerang delegation.

And remember, that's what happens when you hand off work to someone. It looks like it's going somewhere, but then it loops back on you needing clarification, correction, completion, and then it's always at the worst possible time. It's not just annoying, it actively derails your own plans for the work you intended to get done that day.

In the capital A Agile world, we have a term for the type of work that is needed only because something went wrong somewhere else in the system. We call it failure demand. And failure demand is work that isn't creating new value. It's picking up the pieces of some previous attempt at building the thing or doing the job right the first time. The Lean world is

Ep #104: Quality Standards for Law Firms: How to Make Expectations Explicit and Work Predictable [Agile Lawyering Part 4]

even less subtle when it comes to talking about this kind of thing. They just call it waste.

And the rework that results from boomerang delegation is a super common cause of failure demand, or waste, in a system, but it's not the only one. And what most of these forms of failure demand have in common is that they're not really about people doing sloppy work or not caring enough. In fact, I want you to view it as evidence of a lack of clear systemic signals. Signals about what work is needed, why it needs to be done, who is responsible for it, and when it should happen.

And when your delivery team doesn't have good ways of defining those signals, then it becomes prone to a form of signal distortion. It's kind of like when you're on a road trip out west and trying to pick up a baseball game on AM radio. If you're far from the transmitter, or maybe there's a couple of stations broadcasting on the same frequency, it can be hard to figure out what's going on. The signal gets distorted and that distortion causes uncertainty.

In a law practice, signal distortion shows up in a few common patterns. I've already talked about boomerang delegation, and it's definitely one of those patterns. Another one is something I call shotgun delegation. And this is one of those odd but foreseeable results when internal stakeholders, and especially those in positions of power, have anxiety around predictability.

And it happens when someone, it's often a senior lawyer, but sometimes it could be a client or another stakeholder, and they know they need something done, but they either don't understand or they don't trust the team's process for getting that thing done. So instead of sending a clear signal about their need to a single well-defined workflow, they send it everywhere.

Ep #104: Quality Standards for Law Firms: How to Make Expectations Explicit and Work Predictable [Agile Lawyering Part 4]

I actually facilitated a workshop a few months ago where a senior lawyer described this perfectly. By their own admission, they said, quote, "I don't always know how I'm supposed to get people's attention, so I try everything. I'll make a task in Cleo, I'll send an email to the whole paralegal team, and I'll post a message in Teams. I just need to know that someone is getting my request and they're working on it." End quote.

Now, I understand that instinct, but think about what that does to the rest of the team. Are they expected to monitor all those channels? Who's supposed to take the lead? How will they or the attorney know how to prioritize it? And what channel should they use to respond? Ultimately, do they need to tie off the loose ends in each of those systems or does responding in one of them suffice?

So even if that senior attorney's approach is understandable, there's no denying that it is chaotic and expensive in the way that it takes up multiple people's time and capacity. By shotgunning the request across multiple channels, what they're actually doing is increasing the amount of noise in the system overall. And counterintuitively, they're making it less likely that their signal will get picked up and processed correctly.

So that's shotgun delegation. The last form of signal distortion I want to name for now is something I've started to call urgency camouflage. This one also comes from anxiety around predictability. And urgency camouflage happens when someone's need isn't really urgent, but they're worried that if they present it honestly, it's going to fall to the bottom of everyone's to-do list. So they dress up the request in urgency signals in an attempt to make people treat it as high priority.

And sometimes those signals show up as bombastic behaviors, raised voices, dramatic language, or pressure tactics. A lot of times they're more subtle, false deadlines, sometimes given in the hope that the work will

Ep #104: Quality Standards for Law Firms: How to Make Expectations Explicit and Work Predictable [Agile Lawyering Part 4]

actually get done well ahead of the real need. And sometimes they show up as outright subterfuge.

People who are adept at office politics often develop ways to sort of smuggle their requests into the workflow through side channels, maybe tapping a particular associate or paralegal that they know they can persuade to move the work ahead of everything else, get it to jump the queue.

My overall point is that all three of these patterns, boomerang delegation, shotgun delegation, and urgency camouflage, are closely related because they're all forms of failure demand. And they're all driven by distorted or unreliable signals, and they all emerge when people don't trust the system to deliver work consistently and predictably.

Also, by the way, none of them are solved by getting people to behave better or telling them to communicate better or just try harder. They're solved by making expectations explicit, by defining when work is actually ready for someone's attention to move it forward, and then what done looks like so they can confidently stop and move on to the next thing.

In other words, we can solve for nearly all of these problems by creating quality standards that make the signals trustworthy again. So that's where we're headed next. Now, there's a difference between a quality expectation and a quality standard. And in most law practices, quality expectations exist, but they often exist almost exclusively in people's heads. And this tendency to define quality as, "Well, I know it when I see it," leads to two distinct problems.

First, implicit quality standards force mind reading. The person doing the work has to guess what the requester actually wants: what level of detail, what format, what analysis or documentation should be there, what context

Ep #104: Quality Standards for Law Firms: How to Make Expectations Explicit and Work Predictable [Agile Lawyering Part 4]

do they need? And when those expectations aren't explicit, then all of that uncertainty gets pushed onto the person doing the work.

That increases their cognitive load. It makes rework far more likely. It also makes it more likely that they're going to put off getting started because they're hoping that they're going to get a better idea or ideally better instruction around what they're actually supposed to do.

And this is another one of those situations where our human desire to be, quote, "efficient" in the short term actually sabotages the system's efficiency and effectiveness over the long term.

Yes, it seems easier and faster to quickly assign a piece of work to someone else with little context. But if that work comes back and doesn't meet your expectations, your brain will make up any number of stories to justify your decision, why what you did was okay and why the lack of quality is therefore the other person's fault. You might think, "Well, they're supposed to know this as part of their job," or it might say, "They're a smart person, they should be able to figure it out," or maybe, "Well, if they can't do this right, then maybe they're not the right person for this role."

I'm here to tell you, blame is cheap. The reality is, if you're the one who needs the work done well, then it is ultimately your job to create the environment for others to succeed. And the quality standards method I'm going to teach you in a few minutes are a great way to start building that environment. This is the ounce of prevention that avoids that pound of cure later.

The second issue with implicit quality standards is a bit more uncomfortable. Even when those notions of quality live inside a reasonably competent person's head, they aren't nearly as consistent as we might want to believe, because human brains are squishy. Even highly trained

Ep #104: Quality Standards for Law Firms: How to Make Expectations Explicit and Work Predictable [Agile Lawyering Part 4]

professionals are influenced by context, by time pressure, fatigue, interruptions, whatever else is competing for attention that day.

And what that means is what feels like my standard way of doing this actually drifts over time. And once work moves across multiple people, across multiple brains, then that variability compounds.

I'm a big fan of a particular book on this topic, *The Checklist Manifesto* by Atul Gawande. And people often assume that this is a book about productivity, but the subtitle is *Getting Things Right*. This is a book about quality control. And the whole framing of the book is Gawande trying to figure out why a particular surgical procedure had such an unacceptably high complication rate despite surgeons being some of the smartest and most highly trained members of society.

And he isn't picking on surgeons because they were careless or incompetent. He's illustrating that even those smart and educated professionals are still subject to inconsistency and mental lapses. We all are.

But by introducing a very simple tool, the quality control checklist, and then empowering the entire team with the authority to enforce that quality standard, Gawande showed how the rate of complications fell precipitously. Or, put another way, by making quality standards explicit and workable, the surgical teams were able to predictably and consistently achieve successful patient outcomes. And the same logic applies in your law practice.

When we make quality standards explicit, we're not trying to micromanage professionals or turn people into automatons. We're acknowledging reality: that complex work performed by humans across multiple handoffs benefits from shared agreements about what matters and what good enough looks like in context. And as a quick side point, the same thing is true if you're

Ep #104: Quality Standards for Law Firms: How to Make Expectations Explicit and Work Predictable [Agile Lawyering Part 4]

starting to outsource more and more of your work to machines, including AI.

From a flow perspective, explicit quality standards do three critical things. They reduce rework by making expectations clear up front. They reduce duplicate effort by clarifying who's responsible for what. And they reduce the amount of effort needed at that quality review stage by making it easier for the reviewer to trust that by the time work gets to QA, it's likely to pass that review.

All of that improves consistency. It keeps the work flowing towards done, and because the work is flowing more smoothly, it's also one of the main ways we create predictability in the system overall. And importantly, like the improved patient outcomes Gawande writes about, it also improves outcomes for clients. And I'm not going to dwell much on that today, but it's worth saying out loud: Smoother flow isn't the end goal. Better client outcomes are. I'm going to come back to that later in the series.

For now, the takeaway is simple. If quality standards remain stuck in people's heads, then flow will be fragile. Making your policies explicit is one of the most powerful stabilizing moves you can make for trust, for predictability, and for the overall client outcomes like I said.

Okay, so up to this point, I've been talking about quality standards in a pretty abstract way. Now I want to get down to details. The first and most important form of an effective quality standard is what we in the Agile world call a definition of done. And the definition of done is actually pretty straightforward. It needs to answer this question: What are the key things that need to be true or at least accounted for in order for this deliverable to be fit-for-purpose?

Now, notice a few things about that framing. One is that I mentioned a deliverable. And this is something we don't always consider in a law

Ep #104: Quality Standards for Law Firms: How to Make Expectations Explicit and Work Predictable [Agile Lawyering Part 4]

practice environment. We often think we can communicate information to someone and sort of be done with it. And if you do that verbally, that's not really a deliverable. At best, it's sort of an ephemeral one.

When we're talking about quality standards and quality improvement, I'm going to strongly encourage you to think about working towards something that is tangible, because when it's tangible, a document of some sort, you can define what that deliverable is supposed to look like and you make it easier for someone else to review that work against your definition of quality.

Now, your deliverable doesn't have to be fancy, but if you're going to communicate something verbally, then maybe your tangible output of a legal research process is a bullet point outline, if not a full script.

Better yet, how much more durable would your effort be if you captured it in a simple research memo that you can then send to someone or put in the file after you've had your verbal discussion about it? Now, if you're going to send information just over email, then maybe having a template of some sort can go a long way to making that email more effective. And of course, if it's substantive legal work, like a contract provision or a court filing, then you already know how useful and effective templates can be.

My point is that once you have that tangible output, it becomes easier to define what done looks like. Have all the sections in the form or template been filled out? Have the headers been updated? Have you written in an executive summary or TLDR? Has the correct rule or statute been cited and explained? Have specific names or facts been double checked against our records? Are we clear on filing or delivery instructions?

Now, obviously, I'm generalizing a bit, but these are all pretty common needs that so many lawyers I see just assume everyone in their practice

Ep #104: Quality Standards for Law Firms: How to Make Expectations Explicit and Work Predictable [Agile Lawyering Part 4]

should know. And maybe they should, but there's almost no risk in making those assumptions explicit, and there is so much upside.

When definitions of done stay in the realm of, "Well, we talked about it," or, "I told you this," without actually writing them down, then they fall to that squishy human brain problem. And when they're squishy, it becomes really difficult to have a meaningful feedback loop around whether those quality expectations were actually met.

So a strong definition of done starts by clearly describing what that tangible output needs to look like to be fit-for-purpose for the thing. It doesn't need to be overkill. In fact, it shouldn't be. I tell my consulting clients to shoot for like three to eight things at first. Sometimes a complex process or deliverable will eventually need more, but banging out an initial list of three to eight is something most teams can do in like 10 minutes or less.

Now, one other thing to take into consideration when defining a deliverable and what a fit-for-purpose version of that deliverable should look like is to ask an important follow-up question: Who is this output for? Every deliverable has a customer. Sometimes that customer's obvious: a court, an agency, or a client.

But a lot of times legal deliverables have multiple customers. A filing that goes to the court will meet the needs of the judge and opposing counsel, but it might also be read by your client. And each of those audiences has different needs and expectations, and a definition of done that considers one of them is more likely to create some form of failure demand later from the other audiences.

And I don't want to make things too complicated, but creating separate deliverables can be really useful in this situation. And it's one of the genuinely useful applications of AI in legal work that I'm seeing.

Ep #104: Quality Standards for Law Firms: How to Make Expectations Explicit and Work Predictable [Agile Lawyering Part 4]

The court filing can remain the court filing with its own clear definition of done for the judge and opposing counsel, but then you can use AI to create a short explanation or summary letter that helps the client understand what just happened and why it matters. The important point here isn't about AI itself. It's that once you start being explicit about what done looks like for different customers, it becomes easier to see when a single deliverable is sufficient and when it isn't.

Now, I can kind of hear you saying, "All right, that all sounds good, but I don't have time to develop standards and templates for every little thing that I do. That would be so inefficient. I've got work to do." And I get it, but a couple of things.

Number one, this work is an investment in the health of your practice and the sanity of the people who work within it. And my advice is don't try to boil the ocean. For the work that you are doing, take the extra five to 10 minutes to ask in the moment, what does done look like for this thing? What are those three to eight things we need to make sure are true or at least accounted for in order for this deliverable to meet the need at hand? And honestly, if you're still billing hourly, this is probably billable work.

Then capture it, right? Make sure that it's preserved, throw it in a share drive on an intranet, just somewhere so that you are able to get it into a place where it's going to be repeatable the next time that type of deliverable comes around. Now, one last but important thing about a clear definition of done: Most of the examples I've given just now seem to have more to do with making sure things don't get missed, that we don't create something that has insufficient quality. But remember that creating fit-for-purpose work also means not overbuilding your deliverable.

So another important role of a definition of done is that it creates a clear stopping point. We don't want people to stop too early because that leads to waste coming from rework, but we also don't want people to keep

Ep #104: Quality Standards for Law Firms: How to Make Expectations Explicit and Work Predictable [Agile Lawyering Part 4]

polishing and refining that deliverable long past the point of diminishing returns.

So a well-defined definition of done helps prevent both under quality and over quality. It gives people permission to stop, to trust that the work they've done is sufficient for its purpose, and to move on without carrying the cognitive burden of wondering whether they missed something.

So in an important way, definitions of done don't just improve quality, they protect capacity. They also reduce anxiety and ultimately support the consistency and predictability that allows work to flow smoothly in your practice, what I keep saying. Now, just creating that definition of done will help tremendously inside of your systems. But if you're ready to go a little bit further, then once you have a clear definition of done, the next quality standard to think about is the definition of ready.

So if the definition of done asks the question, "what does finished work look like," then the definition of ready answers a different but equally important question: What needs to be true or at least accounted for before work on this deliverable should even begin?

So when I talked about definitions of done, I talked about deliverables and customers, about making sure outputs are fit-for-purpose and meet the needs of the people they're intended to serve. And that way of thinking naturally raises questions about inputs as well.

And in the Lean world, there's a great model that helps make this visible. It's called a SIPOC chart. And SIPOC is an initialism that stands for suppliers, inputs, process, outputs, and customers. And I don't want to turn this into a jargon lesson, but the basic idea is straightforward. Before you can reasonably expect a process to produce a good output, you need the right inputs, and those inputs need to come from somewhere.

Ep #104: Quality Standards for Law Firms: How to Make Expectations Explicit and Work Predictable [Agile Lawyering Part 4]

In legal work, those inputs are very often information and documents. Might be client provided materials, internal research, factual details, instructions, approvals, decisions that need to be made upstream. And a definition of ready is simply a way of making those input requirements explicit too.

And without a clear definition of ready, the risk is that you start working on something that isn't yet ripe for action. People begin working without the things they need or with unclear instructions or without knowing what the problem they're actually trying to solve looks like. And that doesn't just slow things down, it creates rework and frustration and a lot of failure demand later.

So it can be critically important to clearly define what are the inputs you're going to need and where are you going to get them before you actually begin doing the processing part of a particular piece of work.

Now, there's also another type of flow problem that can show up that is related to definitions of ready, which is people doing work that feels like it's a good idea in the moment, but the process is actually designed to account for somewhere downstream. And this happens when people don't understand the full workflow, and when, especially if the quality standards are implicit, they wind up rabbit trailing, doing that extra analysis or drafting or side work now because they're worried that it's not going to happen later.

So in reality, we're looking for a good system that will have clear definitions of ready, clear definitions of done, and that way people who understand the whole thing will already have confidence that there's a place for the work that needs to happen. When people don't have that confidence, they tend to go rogue.

Now, that doesn't mean that we want our definitions of ready to be rigid or bureaucratic, right? Like with definitions of done, they should start simple, a

Ep #104: Quality Standards for Law Firms: How to Make Expectations Explicit and Work Predictable [Agile Lawyering Part 4]

short list, three to eight things that need to be true or accounted for before work begins.

And again, you don't have to go do this for every single piece of work that might happen anywhere in your system. You can start by thinking about this in the context of the work that is on your plate right now. Do I have all the ingredients I need to cook this recipe? Can I get started right now and be confident that I can complete this deliverable in a single sitting without having to go find other things before I can actually get done? Then, if you keep track of it for the thing you're working on now, eventually you'll build up a library of these things that you and your team can use over time.

Also, I can tell you that over time, as teams mature and workflows stabilize, something else interesting is going to happen. And that is that the definition of done for one stage is going to start to merge with the definition of ready for the next stage. That way the quality standards become these shared expectations across the entire system as opposed to sort of checkpoints imposed from the outside.

Now, one other thing I will tell you for both definitions of ready and definitions of done is as much as you can, resist the urge to develop them in a vacuum. If you can have a conversation with the relevant team members who are involved in the process about defining what ready looks like, and then again, what done looks like, that's going to be a way more durable and ultimately better way to come up with a complete standard that works for you and your team.

It is way easier to get buy-in when people are involved in the co-creation of the thing than it is for you to sort of hand it down from on high and then explain to people what it is that you want. It's better if we all work together to develop what we want.

Ep #104: Quality Standards for Law Firms: How to Make Expectations Explicit and Work Predictable [Agile Lawyering Part 4]

Okay, one last point before I wrap up for today. Even with clear definitions of ready and done, quality misses are still going to happen. That's normal. Process improvement is itself a process, and the goal isn't immediate perfection. The goal is to just start with something, learn how it works, and then adjust and improve. And by just getting started and adjusting along the way, then over a surprisingly short amount of time, you're going to wind up with quality standards that are both highly effective and usable by your teams in their day-to-day work.

So let me give you a few quick thoughts about making those adjustments. When some piece of work turns out to fail the quality check, it doesn't meet the quality needs, then there are a couple of things you need to do. One, obviously, is to correct the error or otherwise do your best to make things right. But then after that, it's essential that you and the organization learn from the failure. And, like I said at the top of the show, I'm going to encourage you to resist the urge to chalk it up to a people problem.

Now, that may be the case, but more likely the failure is signaling a process improvement opportunity. In general, when you discover a quality problem and you've used a quality standard, then only one of two things is likely to be true. One is, if the quality standard was followed but the error happened anyway, then the policy needs improvement. Either something is missing from the standard or it wasn't written clearly enough. And so the solution there is to fix the standard so that as new work comes down through your system, it never has that same type of problem again.

Now, if that quality standard exists but it wasn't followed, then there's another set of issues. It might be an education problem if the person either doesn't have the training or the context to properly follow the procedure. It might be a findability problem. Maybe the standard exists but the person didn't know where to reference it or how to find it.

Ep #104: Quality Standards for Law Firms: How to Make Expectations Explicit and Work Predictable [Agile Lawyering Part 4]

If the standard exists but didn't get followed and neither of those other things are true, then maybe you've got a people problem. That's where you've got an engagement or an attitude problem, and you'll need to figure out that's the place where maybe the person isn't the best fit.

But I will keep cautioning you against defaulting to that last one, because it's where a lot of leaders' heads want to go, but I think it's the rarest case. Most quality problems are systems issues, not character flaws.

And if you treat quality failures as opportunities to diagnose a system improvement opportunity rather than default to asking who messed up, then you're going to create the conditions for continuous improvement instead of defensiveness. And that's how quality standards and organizations get better over time, not by adding too many rules, but by refining shared understanding based on real feedback from the work itself.

So let me wrap up and leave you with this. In today's episode, I've been talking about quality as fit-for-purpose consistency, not perfectionism, but as the guardrails for flow. Definitions of done and definitions of ready are simple tools, but together they stabilize flow by reducing rework, preventing backflow, and preventing surprise interruptions. And in the context of the series, that stability matters because without it, it's almost impossible to make credible commitments. Your promise-keeping machine will have friction.

But when the workflow is stable and predictable, you'll see the velocity of your delivery system start to increase. And yes, despite my earlier cautions, I'm talking about speed now, but it's speed that works because the right guardrails are in place. Better yet, once you're able to trust that work will start at the right time and finish cleanly, you'll be in a way better position to understand how much work your system can comfortably manage at once. And that's where I'm going to pick up next week.

Ep #104: Quality Standards for Law Firms: How to Make Expectations Explicit and Work Predictable [Agile Lawyering Part 4]

In the next episode, I'm going to introduce a specific form of explicit policy called a work in progress limit or WIP limit. WIP limits are an amazingly effective tool, but for many people they are counterintuitive at first. What I can tell you though is that by moderating your number of concurrent commitments in the short term, you and your team will be able to deliver way more things over the mid to long term. So stay tuned for that.

So let me briefly connect today's discussion back to how GreenLine supports these ideas in practice. We've designed GreenLine so that quality standards live in the flow of work itself, not in a separate policy library or knowledge base that people have to break their flow to go check. Definitions of done, exit criteria, and other lightweight quality gates are embedded directly into the phases where work actually happens.

We're also careful to avoid a common trap I see in a lot of law practice management tools where they create this avalanche of task assignments and automated, unrealistic deadlines that turn dashboards into this wall of overdue tasks, many of which were never truly ripe for work to begin with. And that kind of noise winds up adding to your anxiety and overwhelm, not relieving it.

And I actually find it a little bit funny, although also a little bit sad, that some of those companies are now trying to sell you an AI package on top of it all, purportedly to help lawyers sort through all that noise to find the right signal of work that actually needs their attention.

And at GreenLine, we think the better approach is to reduce the amount of noise in the first place. And through our hands-on onboarding, we work with you and your team to find that fit-for-purpose balance. And then we help you configure your workflow to stay within that Goldilocks zone where quality standards prevent rework and reduce overwhelm instead of adding to it.

Ep #104: Quality Standards for Law Firms: How to Make Expectations Explicit and Work Predictable [Agile Lawyering Part 4]

To see how this works in practice, you can head on over to greenline.legal to learn more. And you can also book a demo to see GreenLine's approach to embedding quality standards in action.

All right, that's it for today. If you found today's episode useful or enlightening, I would really appreciate it if you would share it with a friend or colleague who could maybe benefit from a more Agile approach in their legal practice. And of course, to keep the rest of this 101 series coming for you, please hit the follow button in your favorite podcast player.

If you have thoughts, questions, or topics you'd like to hear me discuss, please don't hesitate to reach out to me at john.grant@greenline.legal.

As always, this podcast gets production support from the fantastic team at Digital Freedom Productions, and our theme music is "Hello" by Lunareh. Thanks for listening, and I will catch you again next week.