

Ep #122: The AI Productivity Trap: What Lawyers Are Missing About ROI [AI ROI Part 1]



Full Episode Transcript

With Your Host

John E. Grant

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

Ep #122: The AI Productivity Trap: What Lawyers Are Missing About ROI [AI ROI Part 1]

The pitch for using AI in legal goes something like this. The machines are going to handle the grunt work, the tedious stuff, so that you can focus on the high-value stuff. And somewhere in there, you get your evenings and weekends back. It's a compelling pitch, but the data is starting to come in, and it's telling a different story than the one we're being sold.

The benefits of AI can be real, some of them, in some contexts, but they're not as consistent or frictionless as the pitches suggest. And here's what almost nobody is talking about. Even if the benefits are real, what's on the other side of the ROI equation? The investment part. Are we doing an honest accounting around the time, the energy, the quality risk, the effect on your capacity?

Return on investment requires a clear-eyed assessment on both sides of the ledger. Today, I'm going to suggest some better ways to start thinking about them.

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.

Hey everyone, welcome back. So, there is more than enough commentary about AI and legal right now. You would fill a dozen bookcases, 100 podcast feeds, and I've been kind of reluctant to add my voice too much to that conversation, at least until now. But I think it's the right time because so much of what's out there is either breathlessly optimistic or catastrophically pessimistic, and neither of those is particularly useful if you're trying to figure out how to really use this tool to better run your law practice.

And I'd say if I had to position myself somewhere in that spectrum, I'd call myself a cautious pragmatist. But now that we're three-plus years into this

Ep #122: The AI Productivity Trap: What Lawyers Are Missing About ROI [AI ROI Part 1]

large language model era of AI, I think we've reached a point where it's worth taking a step back and thinking about AI not just in terms of what it can do or what it promises it can do, but how should we actually be evaluating it as a business tool? Not so much whether to embrace it or shun it. That debate is, I think, mostly over.

The AI is here. The question, to my mind at least, is how to use it well, how to use it safely, and how to use it in a way that benefits you and your clients.

So, I'm using today's episode to kick off what I'm going to call an occasional series, and I'm naming it AI ROI, return on investment. And I want to be clear from the start, when I say ROI, I mean something a lot more expansive than the subscription fee versus the hour saved. I'm talking about that genuine accounting of what you're putting in and what you're getting back across lots of different dimensions in your practice, and especially the ones that you feel most acutely, the ones that matter most to you.

And I'm going to start today with something of a framework for even how to think about ROI. So, what it actually means in this context, and then I will give you three specific ways that I think the investment side of the equation is getting systematically underestimated by most people.

So, thinking about it in terms of ROI or in some calculations value is an interesting throwback for me because before I was a podcaster, I was a blogger. And the very first blog that I wrote was titled "Legal Value Theory." And the reason I called it that was that a lot of people in legal were talking about delivering value to clients, value this, value that, but nobody was actually defining what value meant. And I think we've got a similar problem right now with return on investment with AI. Everyone's talking about the return, but we're not taking in the full picture.

Ep #122: The AI Productivity Trap: What Lawyers Are Missing About ROI [AI ROI Part 1]

And so, the definition I used back then was simple: value equals benefit minus investment. It's effectively ROI. But the important thing, and the thing that I was really intentional about making clear, is that both sides of that equation involve a lot more than just money. People trade time for money, money for time, money for enjoyment, all the time. Value is something that you have to define subjectively as much as objectively. In fact, it's maybe more subjective than it is objective.

And back when I started that blog, I read a really dense but fascinating book called *Toward an Anthropological Theory of Value* by a guy named David Graeber. And there was a discussion that he had in the first chapter of that book that really stuck with me. And he observed that in many of the world's languages, the word that we use for economic value and the word we use for our personal values is the same word.

And I think there's a reason why that's true, and he would argue there's a reason why that's true, that when we place economic importance on something, what we're really expressing is these are the things that matter personally, professionally, as a community to us. Economic value and societal values are the same thing.

And so, when I apply that to ROI, I'm thinking about return across probably four dimensions. And you could maybe think of them as four currencies, although even that is sort of an economic word. The first currency is the true financial one, right? Revenue, profitability, cost efficiency, it's the obvious part. Another one, though, is cognitive: your attention, your focus, your capacity to do the work well.

A third one is the currency that is your professional reputation, your judgment, your skills, your ethical standing. And the fourth currency is personal, the fulfillment you get from the work itself and the degree to which the work supports the life you actually want to be living and the impact you want to be having on the world.

Ep #122: The AI Productivity Trap: What Lawyers Are Missing About ROI [AI ROI Part 1]

And I'd argue that a genuinely good return on your investment, whether it's on AI or any other part of your business, has to show up across multiple of those four currencies, ideally all of them. So a tool that saves you money but quietly erodes your judgment or that speeds up production but leaves you too depleted to think clearly, or maybe one that subtly encourages you to take on more work than you can comfortably deliver or deliver with quality intact, that's not a good return on investment. That's a tradeoff that you haven't fully accounted for yet.

And like I said, putting aside the risk part, and I will talk about that in a minute, I feel like a lot of the AI conversation in legal right now is really focused on the financial currency. But I want to flip it over and talk more about the investment side of the equation.

And obviously, part of that is the monetary investment, but I think the real investments here are the ones that don't show up on a credit card statement. It's the time it takes to learn a tool well enough to use it safely. It's the mental energy required to oversee and verify the AI's outputs. And that does require real energy, a lot more than most people expect.

There's the quality risk that comes with any workflow change, but especially one this significant. And then there's the subtle effect that it has on your reckoning with capacity, and I will spend some more time on that in a few minutes too, because I think it might be the most underappreciated cost of all.

Almost everyone I've talked to who's using AI successfully in their practice has told me the same thing. Number one, it takes a lot longer than they expected to get it dialed in. And number two, even then, there are still a lot of problems. That time and energy, all of it, it's an investment. It needs to go on that side of the ledger.

One other thing about just the pure financial side before I move on, the subscription fee is not going to stay where it is. We need to recognize that

Ep #122: The AI Productivity Trap: What Lawyers Are Missing About ROI [AI ROI Part 1]

all of the major AI platforms right now are being really heavily subsidized by investors. We're in this market share growth phase, the old Amazon playbook, the Uber playbook, but it's happening this time at this unprecedented scale.

And they seem to be in this death race where they think that there's only one winner that can come and that eventually they're going to take it all. So investors are really pouring a lot of money and artificially keeping costs low. But we're already starting to see prices go up and that's going to continue. So, even the part of the investment that seems negligible right now is going to deserve a second look before long.

Now, one thing I want to point out as someone who is now solidly middle-aged and has lived through a few other tech exuberation cycles in the past is even in the middle of all of this AI noise, the hype, the doom, whatever, it's not actually a new situation. I remember back when I was in my technology career that we were required to read a book called *Good to Great* by Jim Collins back in 2001. This is right in the middle of the dot-com era where everyone was convinced that the internet had changed every rule of business forever.

And I recently dusted off my copy of that book, and there was something that he wrote that really struck me. He said at the time, quote, "The truth is, there's nothing new about being in a new economy. Those who face the invention of electricity, the telephone, the automobile, the radio, or the transistor, did they feel it was any less of a new economy than we feel today? And in each rendition of the new economy, the best leaders have adhered to certain basic principles with rigor and discipline."

And then Collins goes on to say this, quote, "While the practices of engineering continually evolve and change, the laws of physics remain relatively fixed." And that's a framing that really struck me. The practices evolve, the physics don't. And I think that's right for what we're living

Ep #122: The AI Productivity Trap: What Lawyers Are Missing About ROI [AI ROI Part 1]

through right now. The practices around how we deliver legal work are going to evolve. They already are.

But the laws of physics of a well-run law practice or a well-run legal business haven't really changed. Capacity discipline, quality standards, professional judgment, those are the laws of physics. The AI doesn't change them, but it does raise the stakes on getting them right.

And I think what we're seeing and what Collins was recognizing then and is equally true today is that this pattern where new technology arrives, it promises transformation, it scrambles our mental models, and then settles into this world where the fundamental still govern, it's repeated itself many times.

Email was going to change everything, mobile computing was going to change everything, video conferencing was going to change everything. And each time we got busier before we got better. And each time, the people who navigated it well were the ones that held on to their core principles while adapting their practices to the new technology.

So, one of the reasons I wanted to kick off this series now is that we've actually started to get some data around what AI-empowered productivity actually looks like. And I think it's important to ground the conversation in something more than anecdote. And granted, these are just some early studies, but I think they're useful.

The first one comes from a workforce analytics company called ActiveTrack, and they analyzed the digital work activity of over 150,000 workers at more than a thousand employers. It was one of the largest studies of AI's effects on work habits that we have so far. And what they were looking for was simple: what actually changes when people start using AI tools at work?

Ep #122: The AI Productivity Trap: What Lawyers Are Missing About ROI [AI ROI Part 1]

And the results were not what the productivity optimists had predicted. AI didn't actually reduce work, it intensified it. Time spent on email and messaging more than doubled. The use of the business management tools jumped almost 95%. And the time that the workers devoted to focused, uninterrupted work, the kind of deep thinking that actually moves the needle on complex problems, actually fell by 10%.

AI didn't give people more time to do the high-value work, it stole that time managing all of the outputs and the cross conversations and the quality assurance and all of the things that you need to have in place to work with this technology.

The other one comes from some research from the Boston Consulting Group, BCG, on a phenomenon that they were calling AI brain fry. It made for great headlines. And I think it captures something real. They defined it as a mental fatigue that comes from overseeing AI tools beyond your comfortable cognitive capacity. And they interviewed actual workers. They described it as this sort of mental static, like having a dozen browser tabs open in your head simultaneously.

But the data that they produced was the one that was most interesting. Workers experiencing this AI brain fry made major errors 39% more frequently than those who weren't experiencing that problem. Not 5% more errors, not 15%, almost 40% more errors.

And long-time listeners know, I'm a big fan of Cal Newport, and this is indicative of a pattern that Cal Newport has been describing for well over a decade. He calls it the digital productivity paradox. And what he's observed is that easier consistently translates to busier, not more productive. And it happened with email, it happened with mobile phones, and it's happening now with AI, and it's happening faster.

So, let me get specific, and I want to talk about three particular ways that the investment side of the AI ROI equation is being underestimated in most

Ep #122: The AI Productivity Trap: What Lawyers Are Missing About ROI [AI ROI Part 1]

legal practices right now. These aren't the only costs, but they're the ones I see coming up again and again, and they're the ones that I think the return side of the conversation doesn't really account for.

And the first one is something I talk about a lot on this podcast, and it's just as true for AI as it is for any other process improvement. And that's when you work on the wrong part of your overall system. And as you know, I look at law practices through a systems lens, and one of the things that systems thinking teaches us is this thing, it goes by a few names, the Theory of Constraints. I call it bottleneck theory.

It's that in any complex system, the speed at which the whole system can deliver work is determined by the slowest part of that system. That's the bottleneck. And there may be lots of slow parts, but only one of them is the worst, and only improving that one will actually speed up the flow of work through your entire practice.

Now, the flip side of that is improving parts of your practice that aren't the bottleneck won't help. And there's two scenarios where that comes up. The first is if you use technology to speed something up that is downstream of your actual bottleneck. And the reason that isn't helpful is the work still has to make it through the bottleneck before it gets to that part that you just sped up. Speeding up what comes after the bottleneck doesn't improve the flow of the whole system.

But it's the other one that is more insidious. And that's if you speed up something that's upstream of your bottleneck, maybe it's intake, maybe it's drafting. All you're doing is sending more work towards that log jam faster. You're adding cars to an already congested freeway. The traffic jam at the bottleneck is going to get worse, not better.

And so, before you invest time and energy and money in deploying AI or any technology anywhere in your workflow, the first question worth asking is, is this actually the bottleneck? Because if it's not, you're making an

Ep #122: The AI Productivity Trap: What Lawyers Are Missing About ROI [AI ROI Part 1]

investment that won't yield the return you're hoping for, and it's likely to create new problems in the process.

Now, the second investment cost that I think is being seriously underestimated is the work of quality assurance or quality control. And I've been a little fired up about this one lately, and I've been making the point that, in my opinion at least, legal has a relatively immature concept of the role of quality assurance relative to other industries. And that lack of maturity has made us especially vulnerable to the consequences of poor AI practices. The headlines about it are everywhere.

In other parts of the business world, in manufacturing, engineering, medicine, logistics, software development, there are explicit quality assurance functions. Entire professions are built around the question of how you verify that work product meets the required standard before it goes out the door.

But in legal, we've mostly operated on something closer to a I-know-it-when-I-see-it standard. A senior attorney reviews the work, they apply their judgment, and if it feels right, it goes. And that's not nothing. Experienced judgment is valuable, but it's also not a system, it's a person. And when that person is busy or overloaded, then quality review becomes the thing that gets squeezed or certain parts of the process that they think they've got clear in their head just get missed.

And the unfortunate thing about AI output in that context is it has this particular quality that makes the I-know-it-when-I-see-it standard even less reliable than it already was. And I think partly that's because AI is a really good mimic. It produces work product that looks kind of right. It has the format, the structure, the confident tone of maybe good legal work, whether or not it actually is good legal work.

And then, of course, there's the well-documented tendency in humans to lower our guard when something looks authoritative and polished,

Ep #122: The AI Productivity Trap: What Lawyers Are Missing About ROI [AI ROI Part 1]

especially when it comes from machines. We tend to trust machines. And the research calls it automation bias. There's actually an even fancier word for it called epistemological deference. But whatever you call it, the effect is the same. We accept the output with less scrutiny than it deserves.

And then, of course, part of the problem is that the companies and the people behind these tools, they're selling us on speed and quality. So, we kind of take them at their word that it's going to get delivered and once again, we let our guard down.

Because real quality assurance requires real investment in defining what done looks like, in building explicit review steps into your workflow, and then making sure that a human with the right expertise, enough time, and frankly, the brain space to not defer to the judgment of the machine, actually engaging critically with the output before it goes out the door to a client or a court or who knows where.

So, I think it should be obvious by now. I think it's critically important to get the quality control function right to define the processes, to define the standards, to do the training around it, and then actually preserve capacity to do it well. And that investment is real. We need to account for it when we're adopting these tools.

Now, the third investment is one that I actually talked about back in episode 80 of this podcast, but that was over a year ago. So I'm going to repeat myself here. And I think it's the one that connects directly to that four-currency ROI model I laid out earlier, especially the cognitive currency. Because here's what I'm seeing. AI has induced this sort of magical thinking around capacity. And I think it's because AI does make certain things faster to get started with. We've developed a tendency to believe that our capacity has grown.

And when we think that our capacity is bigger, we will then commit to more work. We say yes more easily. We take on more matters, more tasks, more

Ep #122: The AI Productivity Trap: What Lawyers Are Missing About ROI [AI ROI Part 1]

projects because it feels like we have more bandwidth or at least we should have more bandwidth. But the thing we ignore is that the capacity to get something started is not the same as the capacity to see it all the way through to completion.

And while AI is pretty good at lowering the cost of getting something started, it is not so good at lowering the cost of doing the actual work well: the thinking, the judgment, the quality review, the discussion with humans. All of that still takes time and attention, the same time and attention that it always took, maybe more. But also, you now have to oversee the AI output on top of all the other work.

And I think that's the root of what that ActiveTrack data is actually showing us. AI users aren't actually finishing more work. They're starting more work. They're getting busier, but they're not completing as much. I talk a lot on this podcast about Little's Law, which is the relationship between the amount of work you have in progress and how long it takes you to deliver any single piece of work in your system.

And the short version of this, the more active matters you have in your practice at any one time, the longer it's going to take you on average to finish any single one of them. And AI doesn't change that math. In fact, I think what it's doing is throwing more WIP into the system, more in-progress work, and that actually is creating more delay, not speeding things up.

So, once again, it comes back to the honest reckoning with capacity, and AI is making it harder to do that reckoning. When we commit to more than we can deliver with quality intact, we don't just get busier, we also get into the failure demand spiral that I've talked about before. Quality problems start to emerge, clients get frustrated. You then have to respond to those frustrated clients, you have to do rework on the stuff that had errors.

Ep #122: The AI Productivity Trap: What Lawyers Are Missing About ROI [AI ROI Part 1]

And all of a sudden, we're spending more and more time on the non-value-adding, non-productive work, and each problem feeds into the next one. And I'm not saying that AI can't help you gain a little capacity, but what I am saying is that we have to be even more honest about that honest reckoning with what our true capacity really is.

So, before I wrap up, I want to lean back into that cautious pragmatist role that I defined for myself. And really, none of what I'm saying today is like a pure argument against using AI. I use it in my workflow. I have a lot of clients who use it. I've encouraged its smart adoption on this podcast and in my consulting work. The goal here is not to avoid AI. It's to use it with the same intentionality that you would bring to any significant investment in your law practice.

So, a few ways to maybe help break that down. One is before you deploy AI anywhere, be really intentional about the specific problem you're trying to solve, not what AI can help with. That's too vague. You need to know what's your actual bottleneck. What is the specific place where work is getting stuck or where quality is faltering that you're trying to address? And if you can't answer that clearly, then don't go investing in a solution based on hopes and promises. You've got to get clear on what your actual need is.

Number two, if you are going to adopt AI and if you haven't already, start in lower-risk territory. Think about back-office work, policy drafting, marketing content, some administrative templates. It's a way safer place to learn about the tool's capabilities and its failure modes than getting in and doing the substantive client work. It's a lot less risky. And that'll give you a way to understand where and how AI gets things right and gets them wrong before you start applying it into the more consequential client-facing work.

The third one is you've got to be intentional about building quality assurance into your workflow before you need it, not after something goes wrong, that's too late. We want to know what done looks like, right?

Ep #122: The AI Productivity Trap: What Lawyers Are Missing About ROI [AI ROI Part 1]

Definitions of ready, definitions of done. I talk about that all the time. We want to talk about what are the things that have to be true or at least accounted for before we will let this work out of our system.

And I feel like it should go without saying, but every day, there's yet another headline about lawyers that have let actual legal briefs out the door without verifying that the cases they cite say what they mean to say or that they even exist, or that the cases they cite actually support the argument they're trying to make as opposed to some out-of-context quote.

And I should say, these court cases or court filings that are getting people in trouble, they are just the tip of the iceberg. That is one of the fastest feedback loops we have, and I really worry that there are a lot of instances of AI-enabled legal work that are going to come to light over the next years and decades that are getting drafted right now that have exactly the same sorts of hallucination problems or other inventions of the machine that are causing quality problems, but because they don't have opposing counsel or a judge reviewing it in quite the same way, we may not find out about it until there's a problem somewhere down the road, and that can take years.

So, this isn't just a litigator's thing. Whatever your practice area is, you've got to build the right quality standards and then the processes and procedures to apply those standards into your work regardless of what tool set you're using.

And the last one is to just really try to be as honest as you can about your capacity. Don't give way to the magical thinking or the silver bullet stuff that, oh, this is going to make me so much more productive. When you do add AI or any technology to your practice, resist the pull to immediately fill up whatever time it seems to free up with more new commitments. You've got to give yourself time and space to learn what the real efficiency or productivity gain is before you bet your entire practice on it. I talk all the time about people, process, and tools, roughly in that order. You've got to get the first two right.

Ep #122: The AI Productivity Trap: What Lawyers Are Missing About ROI [AI ROI Part 1]

All right, that's it for today. Like I said, I'm going to do other episodes in the series. I'm not sure when the next one will be, but I've got a few in the hopper that I'm working on. But if you've got thoughts or questions about anything in this episode, about how you could be thinking about AI and the return on investment for AI in your own law practice, or even if you've got ideas that you want to share with other lawyers, I would love to hear from you. You can reach me at john.grant@greenline.legal.

And as always, if you found today's episode useful, please share it with a friend or a colleague or give me a rating or review on Apple Podcasts, Spotify, YouTube, wherever you listen.

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.