

Paul-Erik Veel
416-865-2842
pveel@litigate.com

October 5, 2026

On Litigation

In a series of LinkedIn posts (compiled below), I am posting about litigation and legal practice under the banner “On Litigation”. The goal is to build a connected series of short reflections on what it means to be an effective litigator – one post at a time.

Part 1: Introduction

Those who know me know that, in addition to practicing law, I love to write. Over the years, I’ve written full-length journal articles, book chapters, and a steady stream of blog posts. But I’ve always wanted to write a book.

The topic I’ve long envisioned tackling in book form is broad: litigation. When I say “litigation,” I don’t just mean courtroom advocacy, although that’s certainly part of it. There are already excellent texts devoted to oral advocacy. What I have in mind is something broader: a synthesis of the insights, habits, approaches, and decision-making frameworks that great litigators rely on to navigate complex disputes. Put differently, what are the skills, strategies, and tactics that underpin success in complex litigation?

A great litigator, in my view, is someone who can shepherd a dispute from start to finish –maximizing the likelihood of achieving outcomes that matter to the client. I’ve had the privilege of working with and against some exceptional litigators in the course of my work. From observing them, I’ve come to believe that great litigators tend to combine many (if not all) of the following:

- Command of procedure and the rules of evidence
- Deep understanding of the relevant substantive law
- Strong courtroom presence
- Clear, persuasive writing
- The ability to tell a compelling story
- Leadership and team management skills
- Excellent judgment

If I ever do write that book, it won’t be about any one of those skills in isolation. Rather, it would try to weave them together, into a comprehensive picture of what makes a litigator truly effective in high-stakes, complex disputes.

But writing a book will have to wait. Between a busy practice and three young kids, now is not the time for a long-form project.

What I can manage, for now, are LinkedIn posts.

Part 2: What Are We Doing When We Litigate?

Before diving into how we should litigate, it's worth asking a more foundational question: what are we doing when we litigate?

"We try to win cases for our clients" is a common answer, but it's too narrow. Most of what we do happens outside court. "We solve problems for our clients" is broader, but not very useful. That describes every professional.

So, can we define our work in a way that is both accurate and illuminating? I think we can.

Before law school, I studied economics. I later completed a master's degree in economics alongside my JD, and I've taught economics courses, including graduate-level classes on the economic analysis of law. I work closely with economists in my legal practice. So it won't surprise anyone that I think economics offers a helpful lens for thinking about litigation.

From that perspective, here's how I would define what we're doing when we litigate: "We act strategically, under conditions of imperfect information, to optimally advance our client's goals in relation to the expected scope of a dispute."

Let me unpack that. In brief for now, with longer posts to follow.

"Strategic action" means anticipating your opponent's response. Litigation is a game. Not in the sense of being trivial or fun (though it sometimes is), but in the game-theoretic sense: the outcome for each party depends not just on their own actions, but on the actions of others. Your choices shape your opponent's moves, and vice versa. Great litigators think in chains of action and reaction, several steps ahead.

"Imperfect information" is everywhere. Your opponents know things you don't. But more importantly, no one knows exactly how a witness will testify or how a judge will decide. Litigation is human, and humans are unpredictable. Great litigators don't fear uncertainty; they harness it.

"Client goals" can be broader than just winning the case. Sometimes the objective is minimizing cost, protecting reputation, avoiding precedent, or something else entirely. Often, these goals compete. Great litigators know how to weigh those trade-offs.

The "expected scope of the dispute" matters. We're not optimizing for today's motion; we're optimizing for the entire trajectory of the case (or a portfolio of cases). Sometimes that means fighting hard. Sometimes it means retreating. But always with an eye on the bigger picture.

So why does this definition matter?

Because it shows why litigation can be challenging: not because the law is complex (though it can be), but because the decisions are. You're making judgment calls with limited information, against smart opponents, trying to chart the best path for your client not just now, but years down the road.

That's why judgment is so central to great litigation. And while experience builds judgment, so does thoughtful preparation and careful reflection.

Part 3: Strategic Thinking in Litigation

Thinking and acting "strategically" isn't about being clever or theatrical. It's about anticipating your opponent's likely responses to your moves, and shaping your own actions with those responses in mind.

In game theory terms, litigation is an extensive-form game: a branching decision tree where each side's optimal choice depends on how they expect the other side to act later. But an extremely complicated one. Game theory students work with neat diagrams of four or five decision points, a handful of options, and tidy probabilities. Litigation is a different beast. A single case can involve hundreds, sometimes thousands, of decisions, with multiple options at each turn and shifting probabilities as the facts and law evolve.

The clearest example is cross-examination. A skilled cross-examiner (almost) never asks questions blindly. Entire lines of questioning are built around a reasonable expectation of the witness's answers. Every question is informed by the likely range of responses from the witness. That makes cross-examination one of the purest exercises in strategic decision-making in our craft.

But strategy doesn't stop at cross-examination.

Every meaningful decision in litigation – from pleading choices to motions, from discovery requests to trial tactics – will provoke a reaction. The quality of your decisions often depends less on their standalone merits and more on how they will influence what your opponent does next.

Take this example. A plaintiff has a strong $1 million breach of contract claim, but counsel can make a good-faith case for $10

million. The larger claim might pressure the defendant toward a faster settlement and make $1 million seem “reasonable” by comparison.

Or it might backfire.

The defendant could respond by escalating: engaging in a no-stone-unturned approach that uncovers new defences, demanding broader discovery, and pushing the matter into a slower, more expensive litigation track. The defendant may refuse mediation, seeing the plaintiff as unreasonable. And at trial, a judge may be more skeptical of the entire case if the bulk of the damages claim doesn’t hold up.

The lesson isn’t that bigger claims are always good or bad. It’s that there are no universal rules. The right move depends on the context: the parties, the lawyers, the tribunal, and the personalities in play. The best litigators take the time to understand their opponents, their opposing counsel, the decision-maker, and even the key witnesses. The more they understand those actors, the better they can predict reactions and craft actions accordingly.

We can’t know with certainty how an opponent will respond. But uncertainty doesn’t mean ignorance. It means we think in probabilities, not absolutes.

Part 4: Thinking About Probability

I described how strategic thinking in litigation means looking ahead: if we take this step, what will our opponent do next? The challenge is that we never know with certainty how an opponent will respond. What we’re really doing is estimating probabilities of the different things that might happen. So strategic thinking depends on thinking about probability correctly.

But people aren’t always good at thinking about probability. Behavioural economics has catalogued a lot of ways in which our thinking about probability can be skewed. There’s too many to go over here, but here’s a handful that apply in the litigation context:

- The Availability Heuristic. We overestimate the likelihood of outcomes that are vivid or fresh in our minds. A lawyer who just won a trial of a certain type may treat that experience as highly predictive of the next case, even though the facts differ materially.
- Anchoring. Initial numbers or positions exert undue influence. A plaintiff’s inflated damages claim or an aggressive scheduling proposal can shape settlement ranges or procedural timelines, even if the starting point is

unrealistic.

- Overconfidence Bias. Lawyers, like professionals in many fields, tend to overstate the accuracy of their judgments. Predicting a “strong” chance of success on a motion may feel justified, but the reality may be closer to a 50-50 case than counsel’s confidence suggests.

Each of these heuristics distorts how we perceive probability. And when we misperceive probability, we risk misjudging our opponent’s likely moves, a judge’s reaction, or the strength of our own position.

The solution is to recognize our biases and think carefully and systematically about probability. With experience, one’s heuristics may improve, but even experts’ intuitions can be biased. So think carefully and systematically whenever possible.

How do you do that? There’s no substitute for spending the time thinking through the likely outcomes, based on as much information as you can have about the facts, the law, and your opponent.

Speaking for me personally, I find it useful to make explicit probability estimates for critical decision-points, expressing the likelihood of outcomes in percentages rather than vague adjectives. I often model different scenarios with different probability assessments to stress-test how my own assessment impacts the best course of action. Not because the numbers or the scenarios are necessarily “right”, but because it keeps me honest as to how I’m evaluating the best course of action rather than relying purely on gut or heuristics.

However you approach the problem, strategic litigation demands disciplined thinking about probability. Only by confronting risk explicitly can litigators make decisions that reflect both where we want to go and identify the highest likelihood path of getting there.

Part 5: Signaling

Great litigators are strategic thinkers: in deciding what the right move is in any case, they anticipate how their opponent might respond. But strategic thinking goes beyond anticipating an opponent’s response. It also requires thinking about what information your actions convey and, in turn, what that reveals about your strategy.

That implicates the concept of signaling. In economics, signaling describes how actions convey information beyond their immediate effect. Michael Spence, who won the Nobel Prize for his work on signaling in markets, illustrated this through education: a degree is not just about gaining knowledge, but a

"signal" of ability or perseverance to employers.

Litigation is no different. Every move we make signals something to our opponent, and every move they make signals something back. For example:

- Proposing mediation early. This might be read as eagerness to settle, leading the other side to infer you will accept less (or pay more).
- Broad discovery demands. These can signal an intent to pursue no-stone-unturned litigation, indicating a desire to make the process time-consuming and expensive.
- Trial scheduling. Pushing hard for an early trial date can signal eagerness to have the matter decided, suggesting high confidence in the case.

The key is that signaling cuts both ways. We must be alert to what our own actions communicate, since a sound step might inadvertently send our opponent information we'd rather they not have. At the same time, we should read the signals embedded in our opponent's conduct: their procedural choices, timing, and posture often contain valuable clues about how they assess the case, their client's risk tolerance, or their appetite for settlement.

Failing to account for signaling can cause missteps. Consider a party that proposes early mediation as a good-faith effort to resolve matters efficiently. If they overlook that it also signals eagerness to avoid trial, they may enter mediation unprepared for the other side to demand more movement than expected. That doesn't mean you shouldn't pursue early mediation to settle efficiently; it means you must be attuned to the signals your actions convey and take steps without inadvertently communicating what you don't intend.

Litigation is a dialogue. Our actions are part of an ongoing exchange of information, sometimes explicit, sometimes subtle. To be a great litigator, you need to learn to read and convey signals in ways that make that dialogue favourable to your client's position.

Part 6: The Four Things Every Litigator Needs to Know

Cases aren't solved in the abstract. They're litigated in reality, by real people.

Which brings me to the four things every effective litigator should know to win a case:

1. Know the law.
2. Know the facts.

3. Know their opponent.
4. Know their adjudicator.

The first two are obvious to any litigator, so I won't spend more time on those. But many litigators stop at the first two. In fact, the last two are critical to thinking through and executing an effective litigation strategy, because opposing counsel and the adjudicator are the human players through which abstract law and facts are litigated. (Witnesses too, but that's even more complex, and I'll deal with that later.) So let's talk about those last two points a bit more.

Know your judge: You don’t always know in advance who your adjudicator will be, but often you do. And when you do, it would be a mistake not to account for that judge’s background knowledge, their preferences in courtroom style, their history on related issues, and the way they like to process information. The same legal argument can land very differently depending on who is hearing it; points of reference or resonance for one judge may fall flat with another.

Know your opponent: Cases don’t unfold against a blank wall. They unfold against another human being who is making their own strategic calculations. As I've described before, litigation is a strategic interaction: the law and facts may be like an LSAT logic game, but litigation is game theory. If you ignore how your opponent is likely to react, you can build a beautiful legal strategy that collapses the moment it encounters resistance.

Litigation strategy lives in the interplay between these four forms of knowledge. The law and facts are the raw materials. But the judge and the opponent are what determine how those materials can be used, and whether they’ll stand up when tested.

If we stop at the first two, we risk treating litigation like a seminar room debate. But in the courtroom, it’s always real: a real human opponent, and a real human decision-maker. How much time should we, as litigators, really spend preparing for those human dynamics? The best litigators know the answer: more than you think.

Part 7: Balancing Action and Reflection in Litigation

We’ve looked at litigation as a strategic interaction under conditions of uncertainty, explored how probability and signaling shape decision-making, and emphasized the importance of knowing not just the law and facts, but also your opponent and adjudicator. All of this with a view to "act strategically, under conditions of imperfect information, to optimally advance our client’s goals in relation to the expected scope of a dispute".

Together, these posts map out the architecture of litigation strategy. Useful as that framework is, it's also high-level and information-intensive. It helps us understand how to think, but it doesn't always tell us what to do next.

That's where I want to take this series. Going forward, I'll focus less on conceptual structure and more on practical tools, approaches, and strategies for day-to-day litigation.

Before making that shift, one last framework point: there are no universal "rules" in litigation. Every case depends on its circumstances. What works brilliantly in one may be disastrous in another. That's why whatever best practices you adopt, it's essential to build in moments of reflection: times to step back, assess, and recalibrate. The danger is being either too reactive or too relentlessly action-oriented. Good strategy requires deliberately making space to think.

So how do you build that into your practice? A few thoughts:

- Set aside milestone strategy sessions: Go beyond routine check-ins by creating space for broader, big-picture thinking.
- Empower every voice: Junior colleagues may lack experience but often see things fresh. Those perspectives can be invaluable.
- Seek input from people outside of your team: Mentors and colleagues can surface blind spots and sharpen judgment.
- Beware tunnel vision: Don't fall in love with your strategy. Be prepared to pivot.
- Pause under pressure: Some of the best insights emerge in the quiet moments, not the rush.

The best litigators act decisively when needed. But they also carve out time for reflection. Balancing the two is what makes litigators effective.

Part 8: Perfect First Drafts vs. Iterative Work

Litigation is a team sport. No matter how talented the lead counsel, complex disputes are won by teams working in concert. How you run that team is as critical as the courtroom skills you bring to the case. One of the most consequential decisions is how you delegate tasks.

I see two main models.

The first is the perfect first draft: you ask an associate to draft a substantive work product (say, a factum) and expect something "court-ready." (I focus here on legal writing, but the same principles apply to any delegated work product.) The "perfect

first draft" approach demands high effort from the drafter and gives the senior lawyer a nearly finished product. My impression is that this is the dominant model on most litigation teams.

The second is the iterative draft: you ask for an earlier, rough version. It may be skeletal or imperfect, but it arrives sooner and provides a platform to think, react, and workshop together.

Both can work. But I've come to believe the "iterative draft" model often produces better outcomes, and economics helps explain why.

The 80/20 rule teaches that 80% of the value often comes from the first 20% of the effort. A rough draft captures most of the conceptual value – framing issues, surfacing questions – while avoiding the low-return hours spent polishing prose too early. Economists studying innovation likewise show that early, even flawed, prototypes accelerate learning by revealing information you can't predict in advance. Litigation writing is no different: each iteration is a new data point that improves the final product and leaves more time to pivot if strategy needs to change.

There are training benefits, too. An associate asked for a polished draft learns craftsmanship; an associate involved in iterative work also learns judgment: how arguments evolve and trade-offs are weighed.

None of this means the perfect-first-draft model is wrong. Sometimes deadlines or routine work make it the efficient choice. But the key is clarity. These models require different workflows and mindsets. Whether you want perfection up front or iteration over time, be explicit about your expectations. Your team, and your client, will thank you.

Part 9: What Is "Good" Legal Writing, Anyways?

Ask ten litigators what good legal writing is, and you'll get fifteen different answers.

Part of the reason is that people may have different things in mind when they think about good legal writing. It's a mistake to think that there is a single form of good legal writing. There isn't. What counts as good writing depends on both the purpose of the writing and the audience you're writing for.

A pleading or an agreement calls for precision. Every word matters. Definitions are necessary, repetition is to some extent unavoidable, and clarity sometimes requires rigidity. Good writing in that context may be relatively dense and technical.

But don't write your factum like it's a contract. It will fall flat. Factums, or any writing with an advocacy function, require a different kind of clarity: clear structure, persuasive framing, and

simplicity over technicality.

Audience matters too. A reporting letter to an insurer should read differently than a quick update email to a busy general counsel. A factum written for an Ontario judge may look different from one for a BC judge, because the norms and expectations differ. Writing well means meeting your reader where they are.

So if there's no single formula, are there any general rules for good legal writing? I'd say yes, at least for advocacy. Here are four key points to keep in mind:

1. Clarity first. If your reader doesn't understand your point, they won't be persuaded by it. Write as simply as the argument allows.

2. Orient your reader. Use structure, headings, and signposts liberally. Make sure the reader always knows where they are in your argument.

3. Fit the forum. Know the norms and expectations of your audience. A judge distracted by your style won't be focused on your substance.

4. Write to persuade, not to be right in the abstract. Legal writing isn't an academic exercise; it's a practical one. The measure of good writing is whether it moves your audience.

Good legal writing isn't one thing. It's the right thing, for the right reader, at the right time.

Part 10: Treat Every Case Like It's Going to Trial

Most cases settle. Everyone knows that. But the best litigators treat every case as though it's going to trial anyway.

At first glance, that may seem inefficient. Why invest in working up a case like it's going to trial when the overwhelming probability is that it will settle? Economics helps explain why, in most cases, this is the optimal approach to achieving your client's goals.

As I've described in previous posts, litigation involves strategic interaction under uncertainty. Each party acts based on expectations about how the other will behave. Those expectations depend on what each side believes about the other's preparedness and resolve. If you approach a case as though it's destined to settle, you'll signal that in subtle but perceptible ways: limited document work, tentative discovery, half-formed theories. The other side will sense it. And they'll price that into settlement.

The reverse is also true. Litigators who prepare as though they're heading for trial change the negotiation game. They increase the credibility of their threat to proceed, effectively a

commitment device in economic language. In game-theoretic terms, they shift the equilibrium. A lawyer who is truly ready for trial is more likely to achieve a favourable settlement, precisely because they don't need one.

That’s not bravado; it’s expected value. Settlement decisions, like all litigation decisions, turn on probabilities multiplied by payoffs. If you're ready for trial, your expected trial payoff increases (since readiness improves your likelihood of success), and your opponent's expected payoff decreases (since they face a more formidable case). Even if the trial never happens, those revised expectations shape the bargaining range and push resolution closer to your client’s optimal outcome.

There’s another reason to treat every case like it’s going to trial: it disciplines your judgment. Preparing for trial forces clarity. You must decide which facts matter, which legal theories survive scrutiny, and which witnesses you trust. Even in cases that settle early, that trial-focused discipline ensures that every motion, discovery, and negotiation aligns with a coherent endgame.

Treating a case like it will go to trial doesn’t mean acting as if settlement is failure. It means litigating with integrity to the process: developing the record, refining the theory, and making choices that would withstand the light of a courtroom. Ironically, that’s also what makes settlement possible on the best possible terms.

So yes, most cases settle. But the best settlements, and the best advocates, come from those who prepare as though they won’t.

Part 11: Litigating from the Moral High Ground

Every litigator knows that to win, it helps to have the law on your side. But the best litigators know that isn’t enough.

Judges don’t just decide who’s right on the law. They decide who deserves to win.

From the very beginning of a case, great advocates think about both: the legal argument and the moral one. They frame the case so that, by the time the law is applied, the judge wants to apply it in their client’s favour. That doesn’t mean appealing to emotion or abandoning logic. It means showing why the just outcome, morally and institutionally, aligns with your position in law.

Economics helps explain why this matters. Decision-makers, including judges, aren’t neutral processors of information. Behavioural economics teaches that how information is framed affects how it’s perceived. People naturally search for

coherence between what feels fair and what seems correct. When the moral and legal narratives point in the same direction, the decision feels not just permissible, but compelled.

Consider a complex breach of contract case between two sophisticated parties. On paper, it might be a technical dispute over a clause or timing provision. But technical disputes can and should have moral narratives. The plaintiff's narrative might be about reliance and good faith: a party that invested, performed, and trusted, only to be left exposed when the counterparty chose opportunism over obligation. The defendant might instead be about efficiency and certainty: a party who priced risk, complied with the bargain as written, and now faces a claim that would undermine contractual predictability.

Both stories have legal merit. But decision-makers tend to prefer outcomes that feel fair in light of effort, reciprocity, and good faith: all deeply ingrained social heuristics. The side that aligns the formal law with that intuitive sense of justice gives the court psychological permission to decide in their favour.

That's why, when building your case theory, you should think beyond whether you meet the legal test. Ask instead: why should we win?

Because persuasion isn't just about being right in law. It's about making the judge want your side to win, because it feels right, not just reads right. If you can hold both the law and the moral high ground, you don't just argue persuasively. You align incentives, logic, and justice: the conditions under which good judges make confident decisions.

Part 12: The Trial Recipe

Most posts in this series draw on economics or decision theory. This one (mostly) doesn't. Instead, it's something I said to a more junior colleague while we were preparing for trial.

I don't recall what prompted it, but my description of a trial was this: a trial is 25% law, 25% emotional intelligence, 25% poker, 25% theatre, and 50% photocopying.

The law part is obvious. Trials involve proving or rebutting elements of causes of action and defences, within the rules of procedure and evidence. It's foundational. But it's also the part we tend to overemphasize.

Emotional intelligence matters because trials involve real people. And beyond that, trials are stressful – for clients, for witnesses, for opposing counsel, and for your own team. Managing that stress, especially with witnesses, is often the difference between evidence that lands cleanly and evidence that doesn't.

Poker reflects the fact that trials are optimization problems under uncertainty. Every day brings a steady stream of judgment calls: witnesses to call or not, objections to take or let go, points to press or abandon, risks to assume or avoid. Some are planned. Many aren't. You need to exercise judgment under uncertainty – by taking calculated risks – many times a day.

Theatre matters because decision-makers are human. When I say theater, I don't mean acting. I just mean that presentation matters. Trial is about presenting a narrative that is clear, credible, and engaging enough to persuade.

And then there's photocopying. Or, more accurately these days, hyperlinking. The logistics, preparation, and legwork that make everything else possible. Trials always take more work than you expect, particularly on the operational side. Start earlier than you think you need to.

You may have noticed that the math doesn't add up; that's intentional. Trials are intense and always require more than 100%. And that's what makes them both challenging and exhilarating!

Part 13: Why Do Cases Get Tried?

Most cases settle before trial. From an economic perspective, that's exactly what we'd expect.

Settlement avoids legal costs, creating a larger "pie" to split between the parties. It also replaces uncertainty with a known outcome. Economic analysis suggests that if both sides are well-advised and can reasonably assess their chances, almost every case should settle.

And yet, trials still happen. Why? I think the reasons cases don't settle generally fall into five categories.

1. Divergent expectations
The parties genuinely disagree about what will happen at trial. Each side thinks it's more likely to win: sometimes because information gaps persist despite discovery, sometimes because of bad legal advice, and sometimes for reasons that are never clear.

2. Strategic brinkmanship
The parties assess the case similarly, but each believes the other will blink first. Everyone is trying to capture more of the surplus, and no one yields.

3. Extreme outcomes and binding constraints
The structure of the payoffs makes trial rational. A settlement and a loss at trial may be equally bad for a defendant, while only a win is acceptable. Litigation funding can, depending on the

structure, create similar dynamics for plaintiffs.

4. External incentives
The case matters beyond its dollar value: establishing precedent, deterring other claims, or avoiding the signal that settlement might send.

5. Principal – agent problems
The parties may be best off settling, but the decision-maker for one of them may not be. Managers, insurers, or even counsel can face incentives that diverge from the client's interests.

Why does this matter?

Because our system has tools to encourage settlement, but those tools don't work equally well at addressing all impediments. Discovery, mediation, and pre-trials help where the problem is information or bargaining. They are far less effective where the obstacle is constraints, incentives, or agency problems.

The practical takeaway is simple: if the other side won't settle, it's critical to understand why. Sometimes more mediation helps. Sometimes only time changes the calculus. And sometimes cases just have to be tried.

Knowing which situation you're in is a core part of good litigation judgment.

Part 14: Disciplined Cross-Examinations

Since most cases never make it to trial, most cases aren't determined directly by cross-examinations. But that doesn't mean that cross-examination isn't a critical part of a trial lawyer's toolkit. Cases can be won and lost through effective or poor cross-examinations. Being known to be a formidable cross-examiner can sometimes dissuade your opponents from taking a chance on going to trial. And, most importantly, cross-examinations, if they go well, are one of the most exciting and enjoyable parts of our job.

So let's talk about them!

Let me start today with this point: cross-examination tempts excess that you must resist in favour of disciplined restraint.

Consider: you're cross-examining a witness. You hear answers you don't like. Questions come to mind: some pointed, some clever, some mildly damaging. The instinct is to ask them. That instinct often needs to be resisted.

Cross-examination is not a brainstorming exercise. It is an optimization problem under severe constraints: limited time, limited attention from the trier of fact, and asymmetric risk. Every question has a cost. But not every question has a

meaningful expected upside. If a question is not calculated to produce evidence you might actually rely on, don't ask it.

I see this principle violated most often when it comes to credibility questions. At a very high level, cross-examination does only two things. It can be destructive, aimed at undermining the credibility or reliability of a witness's evidence. Or it can be constructive, aimed at extracting admissions that advance your own theory of the case. These are distinct objectives, and many cross-examinations are weakened by blurring them.

If you are not going to argue at the end of the trial that a witness is unreliable or not credible, destructive questions are largely wasted effort. The occasional "drive-by" credibility shot rarely moves the needle. Judges discount noise. Unless you are prepared to argue that the witness's evidence should not be accepted, nibbling at credibility just dissipates attention that could have been spent building something useful. In that situation, discipline means committing to a purely constructive cross.

The reverse is also true. If credibility is the target, half-measures are worse than restraint. Destructive cross-examination only works when it is sustained, coherent, and unavoidable: that is, when the cumulative effect forces the trier of fact to confront whether the evidence can safely be relied upon at all.

Seen this way, cross-examination is less about clever questions and more about strategic commitment. Good cross-examination isn't just about knowing how to execute on goals. It is also knowing what the goal is, and having the discipline to ask nothing else.

Part 15: The Key Inputs to an Effective Cross-Examination

Cross-examination is often treated as something mystical. Some lawyers are said to "have it." Others don't. The great cross-examiner is imagined as part actor, part instinctive predator: quick on their feet, fearless, impossible to wrong-foot.

That story is mostly wrong.

What looks like instinct is usually preparation. What looks like talent is usually structure. And what looks like improvisation is usually someone executing a plan they built long before the witness ever took the stand.

Effective cross-examination rests on three inputs.

First: a theory.

You must know what you are trying to prove when the case is over. That means knowing the legal test and having a coherent factual story that fits it. Every meaningful cross-examination is backward-looking: it is designed from the closing argument outward. If you don't know what you need to be able to say at the end of the case, you will ask questions in the abstract. And abstract cross-examinations rarely help at the end of the day.

Second: a method.

This is the part of cross-examination that most people are taught, and what most people focus on when they think about how to cross-examine. Tight questions. One fact per question. No room to wander. These rules are not sacred, but they exist for a reason: they let you control the narrative and accumulate facts that advance your theory. Good cross-examination is not about cleverness. It is about building something, brick by brick, in a way the judge cannot unsee.

Third: contingency.

Witnesses do not always behave. Answers wander. Defences appear. Evasions creep in. This is where cross-examination is said to become "art."

But this too can and should be prepared.

A strong cross is not a script. It is a decision tree. For every important question, you should know the range of likely answers, and what you will do if you get each of them. When the witness surprises you, you are not improvising from nothing; you are choosing between branches you already mapped out. The best cross-examiners are not fearless. They are, however, rarely surprised.

And that is why preparation, more than temperament or experience, is the true source of what we might label courtroom "instinct."

Part 16: Tweaking the Conventional "Rules" of Cross-Examination

There are well-known maxims that generally make sense to follow, but in my view some of them are misunderstood or applied too rigidly. So I want to propose some tweaks to those "rules".

Rule #1: "Don't ask a question you don't know the answer to."

This advice is frequently given, and it is certainly preferable to know the answer before you ask. But it is not always realistic. A better formulation is that you should never ask a question unless you are prepared to deal with whatever answer you might receive. If you do not know the precise answer, but the

range of plausible answers is limited and each would be acceptable for your position, asking the question may still make sense. What matters is preparation, not omniscience.

Rule #2: “Only ask closed questions.”

This is good advice the vast majority of the time. Closed-ended questions should be the default in cross-examination, particularly for less experienced advocates. They are critical for controlling the witness and securing the admissions you need. Even experienced cross-examiners should rely on them most of the time.

That said, they are not always the optimal tactic. Sometimes you need to understand how a witness tells their story before you can challenge it effectively. In those situations, a carefully chosen open-ended question may be stronger than a closed proposition that the witness simply rejects. Used sparingly and deliberately, a witness’s own words can be powerful evidence.

The key is that open-ended questions should only be asked with a plan for addressing the range of possible answers.

Rule #3: “Don’t ask one question too many.”

This rule is often summarized as “stop at the win.” The concern is that once you have the admission you need, further questions invite the witness to undermine it. That risk is real, and it should always be in the cross-examiner’s mind.

But stopping too early can also be a mistake. An initial admission may be ambiguous, easily undone on re-examination, or unlikely to land with the judge as clearly as it should. In those cases, pushing carefully for a clearer or more durable admission may be justified. What matters is weighing the risks and benefits of the next question in light of the witness, their performance to that point, and your overall objectives.

For each line of cross-examination, it helps to know your limited goal, your reasonable goal, and your stretch goal. With a sophisticated or combative witness, achieving only the limited goal may be prudent. With a more cooperative witness, calculated risks may pay off.

The conventional rules of cross-examination are important. They reflect accumulated experience about what usually works. But they are not immutable laws. The best cross-examiners understand not only the rules themselves, but when thoughtful deviation from them makes sense.

Part 17: Listening to the Witness

I want to widen the lens beyond cross-examination and talk about a rule that applies to every kind of examination. It is a

simple rule: pay attention to the witness and the answer they are giving, rather than rigidly following your script.

At a surface level, this sounds obvious. Of course you should listen to the answer before asking the next question. And yet, in practice, many lawyer do not do this nearly as well as we think.

There are a few predictable reasons that lawyers break this rule.

First, good examinations require preparation. Serious preparation. For most lawyers, that means detailed notes. Sometimes it means something very close to a script, complete with contingencies depending on how the witness answers. After investing that much effort, it is tempting to treat an examination as something fixed, something to be executed rather than adapted.

Second, there is a fear of losing control. Lawyers worry that deviating from the script will cause the examination to unravel, that the careful architecture they built in advance will collapse if they go off plan.

Third, examinations can be nerve-wracking, particularly early in one’s career. Notes can feel like a safety blanket. When things feel uncertain, the instinct is often to retreat into the comfort of the page.

All of that is understandable. But it does not change the basic reality: effective examinations engage with the evidence the witness actually gives, not the evidence you hoped they would give.

That means listening and being prepared to pivot. If a line of questions no longer makes sense, abandon it. If the sequence needs to change, change it. If a new question becomes obvious, ask it.

Sometimes a witness will give an unexpectedly helpful answer. Sometimes they will volunteer a fact you did not anticipate. If you are mechanically following your script, you will miss those moments. Other times, the witness will fail to provide the foundation you need for a later point. If you press on regardless, the point simply will not land.

So how do you actually implement this advice? The simplest technique is visual: keep your eyes on the witness, not your notes.

That does not mean you should not prepare. Quite the opposite. I still script every question I plan to ask. Every one. But when I am actually examining a witness, I look at those notes sparingly, and I almost never follow them verbatim.

The preparation shapes my thinking. It helps me understand where I want to go and why. But it is an input, not a set of marching orders.

Preparation is essential. But preparation won't tell you how the witness will actually answer every question. Only listening can.

Part 18: The Law of Evidence Is Not a Problem for Trial

This week's point sounds obvious: think seriously about the rules of evidence well before you get to a hearing. And yet it's one of those principles many lawyers acknowledge, and then quietly forget about until it's too late.

Part of the problem is structural. The law of evidence sits in an awkward place in civil litigation. Administrative tribunals tend to apply it flexibly. Criminal trials apply it rigorously. Civil proceedings fall somewhere in between.

And the early stages of litigation don't tend to enforce evidentiary discipline. On motions, hearsay is often admissible. At the early stages, parties attach exhibits with little thought to whether they could actually be proved at trial.

Then trial arrives, and you need to be attentive to admissibility issues. It's true that some judges will "let it all in" and sort it out as a matter of weight. But others will not, and you don't get to choose which judge you draw. And appellate courts are usually less forgiving when evidentiary discipline slips.

Because of that, it's important to litigate like the rules of evidence matter. If you don't think about admissibility early, you may discover, too late, that you cannot prove facts you assumed were obvious.

Good litigators ask more than what needs to be proven. They also ask how those facts will be proven. They think carefully about whether documents can be proved without calling a witness and, if so, for what purpose. And they think carefully about what witnesses can prove which facts and, if so, for what purpose.

The same discipline applies defensively. Identify the evidentiary weak points in your opponent's case early, not reactively at trial.

And be consistent in your approach. If a category of evidence is inadmissible for your opponent, it may be inadmissible for you. "What's sauce for the goose is sauce for the gander" is one of the core principles of fair trial practice.

The bottom line: evidence is not a trial-week issue. It shapes the case from the start. Ask not only what you need to prove, but how you are actually going to prove it.

Part 19: The Cost/Benefit Analysis of Objections at Trial

Today, I want to deal with another dimension of that same issue: objections.

Many counsel make objections at trial almost reflexively. Sometimes those objections are persistent: every piece of hearsay, and everything that tends toward lay opinion evidence, is objected to.

As with most procedural rights, the fact that you can make an objection doesn't mean that you necessarily should. Every objection consumes scarce trial capital: judicial attention, counsel credibility, and time. The question is whether the anticipated expected value of the objection is positive or negative.

To be clear, I am not part of the camp that never objects at trial. Many objections have positive expected value. Excluding critical testimony or documents that anchor the opposing theory can change the outcome of the case. And where there is a real issue, it is important to create a clean record for appeal. In other circumstances, objections may be important not because they seek to exclude critical evidence, but because they feed into an overall narrative as to the frailties of your opponent's case. In those cases, it makes sense to make the objection, because the expected return justifies it.

However, many objections do not have that positive return. They shave the edges of testimony without changing the shape of the case. They may teach the judge that you object often and gain little. Over time, judges may discount serial objectors. There is also an opportunity cost to objections: when counsel is mentally preoccupied with policing admissibility, they may miss strategic openings that matter far more than the evidentiary infraction they just corrected. Finally, counsel may sometimes come to regret objecting to certain evidence, as it may make it easier for opposing counsel to successfully make a similar objection in the future.

Many good trial lawyers therefore tend to be disciplined about objections. They internalize both the theory of their case and the rules of evidence so deeply that they can choose not to enforce those rules in particular circumstances. When they do stand to object, it is because doing so meaningfully advances their case. This does not mean tolerating inadmissible evidence out of passivity. It means understanding that evidentiary rules are tools, not obligations. The objective is not to enforce doctrinal purity. It is to win the case (or, sometimes, to lose it in a way that optimally positions the appeal).

How does one do that? Making or not making objections optimally is easier when you have considered the potential

evidentiary issues in advance. Potential objections should be carefully considered, in advance of trial when possible, through the lens of a single question: what is the expected payoff of making this objection?

Part 20: The Credibility Economy of Advocacy

Every litigator knows that their credibility with the court matters. But credibility is often treated as a vague professional virtue: something you either have or you don't. In practice, it is more useful to think about credibility in economic terms. Credibility is a scarce resource that can be invested and spent.

Why is credibility important? Because litigation requires judges to make decisions under conditions of limited time and imperfect information. As a result, advocates constantly ask decision-makers to accept what we say. We ask judges to accept our interpretation of facts, our reading of documents, or our characterization of what a case stands for. No decision-maker can independently verify every assertion in a factum or every characterization of the record. At some level, judges must decide which advocates they trust.

That trust has consequences. An advocate who has demonstrated reliability will find that judges are more receptive to arguments that require inference or judgment. By contrast, an advocate who repeatedly overreaches may find even sound arguments greeted with skepticism. How much a judge trusts your argument without needing to see the underlying documentation is, to a significant extent, a function of your credibility as an advocate.

An advocate's credibility is a lot like an investment portfolio. Every statement you make to the court invests a small amount of your credibility. If the judge accepts that statement, you are repaid with interest: the court becomes more willing to trust the next assertion you make. If the judge rejects your statement, that is a default.

And like any investment portfolio, credibility can change quickly. A credible advocate can lose credibility in a moment by overstating the factual record or mischaracterizing the holding of a case. And once credibility is lost, rebuilding it usually requires time and effort to restore the court's trust.

That is why experienced advocates are careful to maintain their credibility.

That does not mean credible advocates avoid difficult arguments or hard cases. They advance both all the time. But they do not paper over the challenges. Instead of overstating the effect of a particular admission, they explain what the

admission establishes and argue why the court should draw the further inference they propose. Instead of pretending that a case stands squarely for a proposition when it does not, they acknowledge that the authorities are imperfect and explain why their position nonetheless follows from the existing case law and broader legal principles. And when they do inadvertently overstate something, they correct the record quickly and candidly.

Advocates who approach the court in that way tend to be trusted by judges. And in the long run, that trust is one of the most valuable assets an advocate can have.

Part 21: Overfitting to Past Cases

Litigators improve by learning from experience. Every case leaves lessons: what arguments landed, what strategies failed, what decisions paid off (or didn't). Over time, those lessons build judgment. That is how good litigators become great ones.

But there is a risk embedded in that process. We sometimes learn the wrong lessons.

When a case goes well, it is tempting to identify the decisions that "caused" the success. The same is true for losses. We search for patterns.

The problem is that litigation is a high-variance environment. Outcomes are shaped by many factors: facts, law, witnesses, opposing counsel, the decision-maker. And there is randomness in each of those.

What that means is that not every outcome reflects the quality of the decisions that preceded it. And not every pattern is real.

In statistical terms, this is overfitting: identifying patterns in past cases that reflect noise rather than underlying structure. "This worked last time" becomes a rule. But it may be a rule built on coincidence.

Behavioural economics helps explain why we fall into this trap. We often rely on heuristics, or mental shortcuts, that provide quick answers. For example:

- Outcome bias: we judge decisions after the results of those decisions are known, biasing the evaluation of the correctness of the decision
- Availability heuristic: we overweight recent or vivid cases
- Confirmation bias: we notice patterns that confirm our instincts
- Hindsight bias: we treat outcomes as predictable after the fact

Each pushes us toward seeing signal where there may only be noise.

This matters because litigation strategy depends on pattern recognition. If the patterns we extract are unreliable, the strategies we build on them will be as well. Over time, that can lead to systematic error: repeating approaches that succeeded for the wrong reasons, and avoiding approaches that failed despite being sound.

None of this means we should ignore experience. The opposite is true.
The goal is to learn more carefully, to separate signal from noise. Here are a few ways to do that:

- Slow down. Heuristics guide our decision-making more often when we "think fast".
- Focus on decisions, not just outcomes. Ask whether a choice was sound given the information available at the time, not whether it happened to succeed.
- Look for patterns across many cases. A single success or failure rarely proves anything. Repeated observations, across different contexts, are more reliable.
- Interrogate counterfactuals. Would the same approach likely have worked if key facts were different?
- Seek disconfirming evidence. Actively look for cases where your preferred approach did not work. Those cases are often more informative than the ones that confirm your instincts.
- Be explicit about uncertainty. Treat lessons as hypotheses with probabilities, not as fixed rules.

Litigation rewards experience. But it rewards well-processed experience more.

Part 22: Mixed Strategies

In earlier posts, I've described litigation as strategic interaction under uncertainty. That framing comes from game theory: your optimal move depends on what you expect your opponent to do, and vice versa. This post leans a bit further into economics than usual (with apologies). But the concepts are worth it.

There is a natural instinct when thinking strategically to assume that at any given decision point there is a single "right" move. Sometimes that's true. But not always.

In some settings, the optimal strategy is not one action. It is a distribution over actions. Game theory calls this a mixed strategy: instead of always doing X or always doing Y, you do X

some of the time and Y the rest. The intuition is that if your opponent can predict what you will do, they can adjust in a way that makes you worse off. In those situations, being predictable is a liability. The solution is not to “choose better.” It is to be harder to predict.

In litigation, we make many kinds of decisions that recur from case to case: whether to bring a particular kind of motion, the breadth of discovery we seek, how aggressively we move a file forward, etc. Each of these can be analyzed on its own expected value. But analyzing the optimal decision is more complicated than that, because your opponent is also learning about you. Lawyers develop reputations. Their patterns are legible: through past cases, word of mouth, and the countless small signals that accumulate over time.

That information is powerful. If you always take a particular step, your opponent will expect you to take that step and build their case with that in mind. If you never take a particular step, your opponent will know not to worry about that. In either case, advance knowledge about an anticipated strategy is valuable.

In some cases, clarity is inevitable. There may only be one clear play for one side, so the other side can reasonably expect that they’ll make that play. But in other cases, it’s less clear, and there may be multiple paths that have similar strategic merit.

This is where mixed strategies matter. If your opponent cannot predict with confidence what step you will take, they cannot optimize against you. They are forced to hedge. And when they hedge, they may not be as effective. That is the core idea behind mixed strategies: unpredictability, in the right circumstances, has value. Not because randomness is inherently good, but because predictability can be exploited.

This is not an argument for abandoning judgment. You are not flipping a coin before deciding whether to bring a motion. The point is narrower. There are situations where:

- multiple options have similar expected value; and
- your choice will shape your opponent’s expectations going forward.

In those situations, always making the same choice can be suboptimal, even if that choice looks locally “right.” In litigation, unpredictability can itself be a source of advantage.

Part 23: Dealing with Bad Facts

Every case has problems. Sometimes it's a document that reads poorly. Sometimes it's conduct that is inconsistent with your case theory. Sometimes it's a witness who presents poorly. And

in many cases, it's all of the above. But that doesn't mean the case can't be won.

Let's talk about how to deal with bad facts. Often, our instinct is to minimize them. That instinct is understandable. It is also often wrong. Bad facts are not static features of a case. They are inputs into a decision-making process. And how they are handled can matter as much as the facts themselves.

There are, broadly speaking, four approaches to bad facts.

1. Ignore them
In this strategy, you don't acknowledge the fact at all, and if your opponent brings it up, you talk about it as little as possible. Sometimes this works, usually where the fact is peripheral. But it’s a high-variance strategy. If the decision-maker seizes on the issue before you address it, you lose control of the narrative. Worse, you risk appearing evasive. And credibility, once lost, is hard to recover.

2. Contain them
In this approach, you don't proactively raise the bad facts, but you're ready to take steps to limit its impact. Maybe you explain why it's not relevant. Maybe it's a matter of contextualizing. A document that looks damaging in isolation may look very different when placed in the full evidentiary record.

3. Confront them
You raise the issue yourself and deal with it directly. Done well, it signals confidence and removes the tactical advantage of your opponent "revealing" the problem.

4. Embrace them
The last option is to not just tolerate the bad fact, but to embrace it. It doesn't always work, but sometimes you can figure out a way to integrate the bad fact into your case theory. Rather than viewing the bad fact as opposed to the narrative, you reshape the narrative to accommodate the apparently bad fact.

The choice between these approaches depends on the nature of the fact, your theory of the case, and how you expect the decision-maker to react. But here are some principles you should always keep in mind:

- Decide on your approach early. Bad facts don’t improve with time.
- Whichever strategy you choose, be ready to deal with the bad fact somehow. A case that only works if certain facts are ignored is not a stable case.

- Be proportionate. Not every bad fact deserves equal attention.
- Protect your credibility. Implausible denial is worse than acknowledgment.

Finally, don't let a few bad facts persuade you that your case is hopeless. You are not required to present a perfect case. You are required to present a better one. Every trial has good days and bad days, and every case has good facts and bad facts.

Great litigators do not win because their cases are free of difficulty. They win because they know how to handle difficulty when it arises. Bad facts are part of the terrain. The skill lies in navigating them.

Part 24: Optimizing the Theory of the Case

Facts are not processed in isolation by decision-makers. They are understood, and decided, as stories. But as you consider a narrative, a problem emerges: how simple should that story be?

On the one hand, simple stories are easier to understand. And easier-to-understand stories are usually more likely to be accepted. But cases are uncertain. If your theory depends on a single chain of reasoning, where every link must hold, it’s fragile. If one part fails, the whole case can collapse.

Because of the latter concern, many lawyers advance overly complicated theories of their case in order to preserve optionality. But that in turn can lead to a complicated presentation that leaves decision-makers scratching their heads to understand the intricate positions.

So there appears to be a tension: simplicity pushes toward a single, clean narrative, while optionality pushes toward multiple paths to victory. More arguments must mean a more complex and contingent narrative.

But that framing is wrong. The mistake is thinking optionality must exist at the level of the narrative. It doesn’t.

A well-constructed theory separates the story from the paths to victory. The story should be simple. It should explain what happened and why you should win, clearly and coherently.

But within that story, you can have multiple independent ways to succeed:

- If one inference is accepted, you win
- If it’s rejected, you still win on a narrower ground
- If the law is applied broadly, you win
- If it’s applied conservatively, you still win

Same story. Different paths to victory. Done properly, multiple paths to victory don’t complicate the narrative. They reinforce it.

How do you do that? Here are a few different ways to think about integrating the optionality into a simple framework:

1. One story, multiple outcomes You don't want to create optionality through shifting or alternative narratives. You want one coherent account with different routes to success.

2. Optionality should be embedded Fallback arguments should feel like natural extensions, not add-ons.

3. Independence matters If every argument depends on the same fact, you don’t have optionality, you have unnecessary redundancy.

4. The goal is resilience A good theory doesn’t just work when everything goes right. It survives when parts of it fail.

The objective is not to choose between simplicity and optionality. It is to build a case that has both: a single, clear story, with multiple ways to win inside it.

Part 25: To Cross-Examine, or Not Cross-Examine: That is the First Question

I have done several posts on cross-examination, but today I want to write about something more fundamental: how to decide whether you should cross-examine a witness or not.

There is an instinct among many lawyers that every witness called by the other side must be crossed. That is wrong. Cross-examination is not an obligation. It is a choice. And like any choice in litigation, it should be made on expected value.

The mere fact that a witness testifies or swears an affidavit for the other side does not justify cross-examination. Evidence exists on a spectrum. Some of it matters to your theory. Some of it doesn’t. Some of it helps you more than it hurts. Treating all adverse evidence as requiring a response is how records become bloated and unfocused.

How do you decide whether to cross-examine a witness? In my view, there are two steps to the analysis: the first is a hard rule, and the second a flexible balancing test.

At the first step, if an opposing witness gives evidence that, if accepted, is fatal or near-fatal to a position that you must win, you cannot ignore it. To be clear, I don't mean that the evidence is harmful; you don't actually need to cross-examine on all evidence that is harmful. For example, if a witness testifies to a fact that only circumstantially supports your opponent's position,

you can often explain it away without challenging the evidence itself. Rather, what I mean is that if the evidence itself is a complete show-stopper for whatever position you are asserting, you must cross-examine. If you leave it untouched, the other side will emphasize that their evidence was uncontested. That is a powerful rhetorical and legal point. You have to avoid that if at all possible.

At the second step, assuming you don't have to cross-examine, it becomes a cost/benefit analysis of whether cross-examination will strengthen or weaken your position.

On the benefit side, there are two distinct objectives.

1. Constructive: Cross-examination can be used to advance your own theory of the case: to obtain facts that fit your narrative, to clarify ambiguities in your favour, or to lock in propositions that you will rely on later. This is the disciplined version of “telling your story through the other side’s witness.”
2. Destructive. Cross-examination can undermine the witness’s evidence, by exposing inconsistencies, narrowing what they actually know, or weakening the inferences the other side asks the court to draw.

On the risk side, the most important risk is expansion of the record in a way that is bad for you. If the witness simply repeats what is already in their affidavit or examination-in-chief, the downside of crossing them is relatively limited. But that is not always what happens. Some witnesses, when engaged, will fill gaps. They will add context that makes their evidence more coherent. They will supply explanations that were missing. They will introduce new justifications that blunt your theory of the case. In doing so, they can materially improve the other side’s record, sometimes in ways that cannot easily be undone.

That cost/benefit analysis requires more than abstract logical reasoning. It requires a view on the witness. What do they actually know? How disciplined are they? How careful? How prepared? Are they likely to stay within the four corners of their evidence, or to expand beyond it when given the opportunity? The answers to those questions will usually be central to your decision.

The key point here is at the first stage, you're asking a question solely about the role the witness's evidence in chief plays in the case. But at the second stage, that question is only one of many, and it's not even the most important one. The question is whether you will go forward or backward by cross-examining the witness, and that usually depends as much on the particulars of the witness as it does on their evidence.

Part 26: Advocacy as a Compression Algorithm

In the late 1990s and early 2000s, early MP3 compression involved a tradeoff. Compress an audio file, and the audio quality degrades noticeably. Songs sounded thin and distorted. But don't compress it, and the song will take hours to download on your dial-up modem.

Over time, the compression algorithms improved. File sizes shrank materially, while the listening experience became increasingly faithful to the original recording. The technology became better at identifying what mattered, what listeners would actually perceive, and what information could be discarded without materially changing the experience.

Effective advocacy in litigation has some real similarities.

Every litigator operates under severe compression constraints. No advocate has as much time as they would like. No judge has as much attention as the case theoretically demands. No factum can include every potentially relevant fact, every passage from every authority, every conceivable argument, or every nuance that exists in counsel’s own understanding of the case.

The advocate’s job is therefore not merely to reproduce every aspect of the case. It is to compress it effectively. The advocate’s internal understanding of the case is the original recording (the .WAV file, if you will). Advocacy is the compression algorithm that attempts to reproduce that understanding in the mind of the judge using dramatically fewer inputs: fewer facts, fewer authorities, fewer words, and less time.

The difficult part is that compression inevitably involves loss. Some arguments do not survive translation well. Some factual nuances require too much context to justify the cognitive bandwidth they consume. Some legal points are technically available but dilute the clarity of the overall signal. Experienced litigators often abandon arguments not because the arguments are wrong, but because they do not compress efficiently.

In some instances, the translation problem is so severe that an argument should not be advanced at all. To stretch the analogy slightly further, a bad argument that will never work at trial is the equivalent of trying to listen to John Cage’s 4?33? on your portable MP3 player. The entire point of the composition depends on context and environment. The problem isn't just one of compression; it's that the medium doesn't make sense. Some litigation arguments function similarly: they may be theoretically sound, but they cannot survive the translation process necessary for persuasive advocacy.

Good advocacy therefore requires thoughtful prioritization. The objective is not to say everything. It is to preserve fidelity to your winning theory. A well-compressed case allows the judge to understand the dispute in substantially the same way counsel understands it, despite receiving only a tiny fraction of the underlying informational inputs.

So what does this mean practically? A few thoughts:

- If a point requires fifteen premises before it becomes persuasive, it is probably fragile when conveyed to the decision-maker.
- If removing an argument improves the coherence of the remaining case, the argument was likely imposing cognitive costs greater than its value.
- If the judge cannot explain your theory of the case back to herself after the hearing in relatively simple terms, the compression likely failed.
- The strongest factums are often not those containing the most information, but those with the highest signal-to-noise ratio.
- Trial and appellate advocacy involve different compression problems. Trials involve compressing vast factual records into coherent narratives. Appeals involve compressing doctrinal complexity into administrable legal principles.

The key thing that combines all of these points is this: advocacy is an exercise in constrained communication. Don't try to transmit the maximum amount of information; try to convey the maximum amount of understanding within the constraints you face.

Part 27: The Value of Open-Ended Discovery

Many litigators have a tendency to approach examinations for discovery as a kind of pre-trial cross-examination. On that view, the goal is to control the witness, ask closed questions, and extract admissions that can later be used at trial or on a motion. There is obvious value in that: admissions matter. A well-planned discovery should identify the propositions you need, and where possible, lock them in.

But that is only part of what examinations for discovery are for. If you approach discovery only as an exercise in extracting admissions, you will get some useful answers, but you will probably also leave without fully understanding the case you are actually facing. In many cases, discovery is where you learn the terrain. And you often do that best not through narrow, closed

questions, but through comprehensive open-ended ones.

That feels counterintuitive to many litigators. Trial Advocacy 101 teaches litigators not to ask open-ended questions during a cross-examination of an opposing witness. That advice is usually right, during a cross-examination. But examinations for discovery are different. The objective is not only to prove your case through admissions. It is also to understand the case, including the narrative the other side will present at trial.

Open-ended questions are valuable because they reveal how the other side understands the dispute. They show what the witness thinks matters. They may identify documents you have not yet appreciated. They reveal who else was involved. They uncover the internal logic of the opposing narrative. All of that matters because the questions you know to ask before discovery are limited by the information you already have. If you only ask those questions, you may confirm or undermine your existing theory, but you may miss the theory the other side is actually going to advance.

There is also strategic value in letting a witness talk. A witness explaining events in their own words usually reveals more than a witness responding yes or no to counsel's propositions. Sometimes the answer is helpful because it contains an admission. Sometimes it is helpful because it exposes confusion. Sometimes it is helpful because it gives you a roadmap for the next round of document requests, undertakings, follow-up questions, or eventual cross-examination.

None of this means discovery should be unstructured. "Tell me everything" is not a strategy. Open-ended discovery works best when it is disciplined. The best approach is usually to move between modes. Start broad enough to understand the witness's account. Then, where it matters, translate the evidence into precise admissions that can be used later. The best discoveries do both. They secure the admissions you need, but they also create space for the witness to explain, expand, and reveal. In many cases, that broader understanding is what allows you to ask the right closed questions later.

Part 28: Know Your Forum

Once upon a time, a litigator would spend most of their time in the same forum. A lawyer practicing in Toronto would practice mostly in the same court in Toronto. Lawyers weren't flying all acr oss the country to different courts, and most lawyers didn't practice in that many administrative tribunals because, well, they didn't exist.

But that's not the reality today. The administrative state has boomed. And it's much more common for lawyers to practice in different jurisdictions. I'm called in Ontario and BC, and I currently have active cases in Ontario, BC, Saskatchewan, and Nova Scotia. I've also done cases in the past in Alberta, Manitoba, and New Brunswick.

A tendency that many lawyers have is to default to the style of their home forum. They default to using the vocabulary of their home forum, the style of advocacy of their home forum, the rules of evidence and objections of their home forum, and the formatting of their home forum. Sometimes that comes from a place of arrogance, but more often it just comes from defaulting to the familiar. Don't do that.

I am firmly of the belief that advocacy is a skill set all of its own and that a skilled advocate can be effective in many different forums. But that doesn't mean it's easy. It requires adapting your advocacy to the local norms, rather than the other way around. It's critical to effective advocacy in whatever forum you're in to demonstrate that you've taken time to learn the local language and rules. That can be hard work, but it is mission critical. Without it, the adjudicator will be distracted by your style and vocabulary at best, and they'll write you off as a dabbler or out-of-towner at worst.

So what do you need to think about? Leaving aside obvious differences in substantive law, here's a list of key things:

1. What do you call the adjudicator? In some forums, "your honour" is respectful. In others, it's inappropriate.
2. How do you refer to basic procedural steps? Is it a motion or an application? Is it a Statement of Claim or Notice of Civil Claim? Understand the rules of court, and use the right nomenclature.
3. How should the material be organized? What are you bringing hard copy? How should things be filed? What's the norm around how the judge reviews the materials? What are the norms around compendiums or books of authorities?
4. What level of detail should your advocacy have? In some forums, adjudicators have time before the hearing to read everything and they only want a brief summary in oral argument. In other forums, good oral advocacy requires you to methodically walk through your entire argument.
5. How do the rules of evidence apply? Different forums take different approaches to all manner of evidentiary issues.

All of those demonstrate humility, by showing that you do what it takes to speak to the adjudicator on their terms. As all great litigators know, that's critical to persuading your judge.

Part 29: Don’t Drink Your Own Kool-Aid

As litigators, we often come not only to identify with our clients, but also to genuinely believe in their cases. At one level, there are good reasons for that. Successful advocacy requires confidence. Clients need to feel that their lawyer understands the case, sees the path forward, and can present it forcefully. A lawyer who seems uncertain about everything is not much comfort to a client facing uncertainty. But there is a difference between advocating a case theory and believing it too much.

Every case theory is a simplification. A case theory takes messy facts, imperfect documents, human motivations, legal uncertainty, and procedural constraints, and turns them into a simplified and coherent whole. That is important. Without a theory, litigation becomes a pile of disconnected facts and issues. The danger is that once we build a theory, we start to see the case through the lens of that theory.

Behavioural economists would recognize the problems. Confirmation bias pushes us to notice facts that support our theory and discount facts that undermine it. The endowment effect can make our own arguments seem stronger because they are ours. Motivated reasoning can quietly turn advocacy into self-persuasion. In short, a lot of different factors make us fall in love with our case.

But when lawyers drink the Kool-Aid on their own case theory, they do a disservice to the client. They stop seeing risk clearly. They become too dismissive of the other side’s arguments. They overread helpful evidence and underread harmful evidence. They prepare for the case they want to have, not the case they actually have.

The client does not need a lawyer who is internally convinced that everything will work out. The client needs a lawyer who can advocate externally with confidence while assessing internally with discipline. Those are different skills, and good litigators need both.

So how do we avoid drinking the Kool-Aid?

First, be deliberate about separating your case theory from your case assessment. They should inform each other, but they are not the same thing.

Second, look for disconfirming evidence. What fact would worry you most if you were on the other side? What argument would you make against yourself? What document would you least

want the court to focus on?

Third, force yourself to state the other side's case fairly. Not as a straw man, but as the best version of what competent opposing counsel will actually argue.

Fourth, update your theory regularly. New information should change your assessment, not simply be absorbed into your initial view.

Fifth, consult your colleagues who will give you honest second opinions.

When you do all of those things, you must absolutely project confidence to both your client and to the court. But the skilled trial lawyer doesn't need to have drunk the Kool-Aid to project confidence.

Part 30: How Many Arguments Are Too Many?

One of the recurring judgment calls in litigation is how many arguments to make.

Lawyers love lists. Some will tell you that you should always have three arguments. That can be useful advice. It is also a rule that is often broken to good effect. I was recently on a summary judgment motion where counsel for another party made 14 points. I have personally made 10 points in a written brief.

So how many arguments should a good litigator make?

There is no universal answer. But there is a useful way to think about it: make the minimum number of arguments necessary to maximize your client's chance of success. That may sound circular, but it's not. It is a practical framework.

Decision-makers have limited bandwidth. That does not mean judges are inattentive. It means they are human. Complexity has a cost. Too many arguments can make a case harder to understand. Weak arguments can crowd out strong ones: if a party advances nine weak arguments and one strong argument, the decision-maker may see a weak case with one better point, rather than a strong case supported by a focused theory.

This is why consolidation matters. Two arguments that depend on the same factual premise and the same legal proposition are often not really two arguments. If they will necessarily rise or fall together, separating them may add complexity without adding value.

But the answer is not to make as few arguments as possible. Often a case has multiple paths to victory. Multiple paths create optionality, hedge against uncertainty, and appeal differently to different decision-makers. The discipline is to distinguish

between arguments that add a new path to victory and arguments that merely restate the same path in different words.

A useful test comes from looking at the incremental value of your argument. List your possible arguments from strongest to weakest. Start with the strongest. Then, going down that list, for each additional argument ask: what does this add, given the arguments I am already making?

The question is not whether the argument has some merit. Many arguments have some merit. The question is whether it materially increases the chance of success enough to justify the added complexity.

Because each argument has a cost. It adds length. It adds cognitive load. It may dilute the theme.

That does not mean the tenth argument is wrong. Sometimes the tenth argument is necessary. But by then, it had better be doing real work. Good advocacy is not an exercise in saying every true thing. It is an exercise in selecting the most persuasive things.

The goal is not to have the longest list of arguments. The goal is to give the decision-maker the clearest path to decide for your client, with enough alternative paths to account for uncertainty, but not so many that the route becomes obscured.

Part 31: AI and the Future of Litigation, Part 1

I was skeptical about writing this post. Everyone is writing about AI and the legal profession. The commentary ranges from over-the-top (“AI will replace lawyers”) to dismissive (“nothing will really change”). The full spectrum was covered; I wasn’t sure I had anything to add.

But that’s never stopped me before. So let me try to apply the same framework I’ve used throughout this series and see if the insights are useful.

Let’s start with what AI is increasingly good at. AI models are, at their core, pattern-recognition engines trained on enormous bodies of text. That makes them genuinely capable in domains where the task involves identifying, synthesizing, or generating text based on patterns in existing material: summarizing documents, drafting first versions of standard-form materials, identifying relevant authorities across a large body of case law, finding inconsistencies in a document set. They’re not perfect, and they still need lawyers to review the output carefully. But they can add a lot of value in certain types of work.

In economic terms, AI is a technology that substantially reduces the marginal cost of certain kinds of cognitive labour. Tasks that

used to take a junior associate hours may take minutes.

By contrast, AI is currently much weaker in domains requiring original judgment under genuine uncertainty, where the right answer cannot be derived from patterns in prior data. Predicting how a particular judge will respond to a novel argument. Deciding whether to bring a particular procedural motion and what effect it will have on the trajectory of a case as a whole. Advising a client on whether to sue at all.

That distinction – low-cost pattern recognition versus high-cost human judgment – is the key economic axis around which the impact of AI on litigation will turn.

The economist William Baumol identified a phenomenon that has come to be known as “cost disease.” Some industries are amenable to productivity improvements through technology. Others are not, because the labour itself is the product. His famous example was the string quartet: no matter how much technology advances, it still takes four musicians the same amount of time to play a Beethoven quartet as it did in 1827.

Litigation is a hybrid. Parts of it look like manufacturing: document review, legal research, drafting standard materials. AI will improve productivity dramatically in those areas. Parts of it look like the string quartet: the skilled oral argument before a skeptical appellate court, the cross-examination of a hostile witness, the judgment call about whether to settle on a Friday afternoon. No AI will replicate those any time soon.

What that means economically is that the value of different kinds of legal work will shift. Work amenable to automation will be further commoditized. The premium will migrate to work that is not.

What does that mean for how AI will impact training junior lawyers and litigation strategy? How does a great litigator use AI to their advantage? More on that next week.

Part 32: AI and the Future of Litigation, Part 2

Here, I want to look at three implications of that reality.

(1) The Junior Lawyer Pipeline Problem

Junior lawyers have traditionally learned by doing. They research, draft, review documents, prepare outlines. That work is not glamorous, but it builds the foundational pattern recognition that eventually becomes judgment. A senior litigator who knows at a glance that a case is weaker than the client thinks has that intuition because they spent years doing exactly the tasks that AI can now perform.

If AI eliminates those tasks, the profession faces a genuine pipeline problem. How do junior lawyers develop judgment if the work that builds judgment is automated away?
The answer isn't to resist AI adoption; that would be like resisting the spreadsheet to train better accountants. But it does mean that legal training will have to evolve deliberately: structured mentorship, simulation, and reflection on outcomes, in place of the accumulated experience that used to come from the work itself.

(2) The Credibility Problem

I described in a previous post how credibility is the litigator's most valuable asset. AI creates new and serious risks to it.

The problem is not that AI is wrong all of the time or even most of the time. It is that AI makes plausible-sounding errors with confident presentation. A factum that cites a case that does not exist, or misconstrues the holding of one that does, can destroy an advocate's credibility before a court, regardless of whether it was generated by AI or written carelessly.

This creates, in substance, a new professional obligation on those who use AI: the duty to verify. AI-assisted work product must be reviewed with genuine critical attention. That is a skill in its own right: knowing what to check, how to check it, and how to recognize the specific failure modes of AI-generated legal reasoning.

This isn't really different in kind from senior lawyers reviewing the work of students. But the ease of creating work product using AI and the confidence of that work product does change the landscape.

(3) Strategic Interaction in AI-assisted Litigation

Return to the definition I offered early on in this series, litigation involves acting strategically, under conditions of imperfect information, to optimally advance our client's goals.

AI changes the information environment. If both sides can rapidly synthesize large document sets, information asymmetries narrow.

It also affects strategic dynamics. If AI tools become widespread, they become table stakes rather than competitive advantage. And in some settings, such as where unpredictability has value, over-reliance on AI-assisted decision-making could itself become a strategic liability.

The Honest Conclusion

AI will change litigation. The honest answer to how much, and in exactly what ways, is that nobody knows. We're all just making

educated guesses.

What I am more confident about: the skills that have always mattered most in complex litigation will continue to matter. Strategic judgment under uncertainty. The ability to read people. The credibility that comes from intellectual honesty before the courts. The capacity to compress a complex case into a simple and compelling story.

Those things have not been automated. I don't think they will be soon.

What AI will do is lower the cost of the inputs that feed into those skills, and raise the value of the judgment that translates those inputs into outcomes. In economic terms, the complementary goods to AI – senior judgment, strategic sophistication, credibility – will become more valuable as AI becomes cheaper.

Part 33: The Joys of Examinations in Chief

I've written a lot in this series about cross-examination. There's no question that it's one of the most exciting parts of a trial lawyer's job.

But most of the evidence a judge actually relies on comes from examination in chief: your own witnesses, telling your own story. And yet chief is often treated as the less exciting part of the trial. That's a mistake. Cross chips away at the other side; chief builds the foundation for everything you argue at the end.

Unlike crosses, examinations in chief aren't about control. They're about narrative clarity. Here are some key points for effective examinations in chief.

(1) Remember you know the case better than the judge. You've lived with this file for months. The judge is meeting it cold. Facts that feel second nature to you—who the parties are, how the industry works, a term of art, the basic chronology—are brand new to them. The instinct is to skip the basics because they feel obvious. Resist it. Build from the ground up, because a judge who doesn't understand the setting can't grasp the significance of the specifics.

(2) Organize so the judge can follow. A witness lived events in some order, but that isn't always the order that makes sense to someone hearing them for the first time. Chronology is a good default, but not the only one. The point is that you impose the structure — decide the shape in advance and lead the witness through it.

(3) Use headings out loud. You can't hand the judge a heading in oral evidence, but you can say one: "Let's turn to the March

meeting." These signposts tell the judge where they are and what's coming, so each answer lands in the right place.

(4) Let the witness talk. The evidence is the witness's, not yours, and its persuasive force comes from their words, in their voice. The anxious lawyer packs so much into a question that the witness just says "yes". Even leaving aside that that's improper leading, it's also not persuasive.

(5) Use documents as a foil, not a checklist. The temptation is to march the witness through the record: "Is this your email? Yes. Your signature? Yes." Those questions authenticate the document but persuade no one. Instead, use it as a springboard: what prompted it, what a line meant, what happened next. The document anchors the evidence; the witness's explanation is what actually persuades.

(6) Know why you're asking every question. Know what each answer needs to establish, so that if you don't get it, you notice and follow up. Work mechanically down a list and you'll pass over the gap. The record is what you argue from, and a fact the witness knows but never quite states is not evidence.

Examinations in chief lack the glamour of cross. But the quiet discipline of building a clear, well-organized narrative is often what wins the case long before closing.

Part 34: Litigating with the Appeal in Mind

Most litigators think about appeals after judgment. That is too late. Appellate strategy runs through the whole case. If I win, will my win survive appeal? If I lose, will I have anything to work with?

To put it in game-theoretic terms, a trial judgment is not the end of the decision tree. It is an intermediate node. Optimizing for the trial outcome alone is solving a subgame in isolation.

There are two conceptual points you need to think about. First, not all wins are the same. The reason is the standard of review. Findings of fact attract deference, while questions of law are reviewed for correctness. An identical outcome at trial can carry a very different expected value depending on how it was reached. A win grounded in the judge's assessment of the evidence is relatively durable. The same win resting on a contestable legal basis is fragile.

Second, the record. If you might lose, you need one: objections taken, positions put, evidentiary or other mid-trial rulings expressly sought. But preservation isn't free. Every objection spends scarce capital—judicial attention, credibility, time—and before a sympathetic judge, aggressive record-building can cost you the very advantage you are trying to protect.

The way to resolve that tension is to treat record preservation as what it is: the purchase of an option. You pay a certain premium now for a contingent payoff: the ability to raise the point later, but only in the state of the world where you lose and the point matters. Its value is roughly the probability you lose, times the probability the point is material on appeal, times the value of the appellate outcome.

So what does that way of thinking mean in practice? A few things:

1. Ask early what the appeal looks like from both sides. Not just whether you would appeal a loss, but what your opponent would appeal if you win.
2. Build wins on findings of fact wherever you can. Deference is an asset. Ask for the findings you need.
3. Prefer multiple independent grounds. A judgment resting on three of them is far harder to overturn than one resting on the strongest single ground.
4. "Price" your record preservation. Where you are likely to win and the point is peripheral, let it go. Where the case is close or the stakes binary, pay a real premium. Where the point will decide the appeal, preserve it regardless; under-insuring a bet-the-company case to keep a judge comfortable is a poor trade.
5. Discount judgments for durability. A large but vulnerable judgment is worth less than its face value once you account for reversal risk, delay, and cost. By contrast, a preserved appellate point is an asset even after a loss at trial. Both sides should be pricing that before judgment, not after.

Trial and appeal are often treated as separate disciplines. They are not separate strategic problems. Litigating as though a trial judgment is the last word is a good way to win something you can't keep.

Part 35: Thinking About the Supreme Court From the Start

Last week, I argued that appellate strategy is not a post-judgment activity: you should be thinking about the appeal from the beginning of the case. This week, one step further: thinking about the Supreme Court of Canada throughout.

This isn't the right strategy in most cases. The Court takes very few, and that low base rate is the right starting point for most files.

But not every case is average. Some carry a genuine legal problem: inconsistent appellate authority, an unresolved question percolating through the lower courts, an old doctrine under new pressure. You can often identify those cases early,

and that recognition should change how you frame your arguments throughout.

It can also change what a "win" looks like. In an ordinary appeal, you are asking a panel to apply established law correctly. The last thing you want is a member of the panel who sees the question differently. The envelope-pushing case is different. There, you are not only playing for the win. If you can't get the win, you should play for a dissent.

That sounds like a strange thing to want, but it isn't. A narrowly framed argument gives a panel nothing to disagree about; naming the doctrinal tension explicitly gives a judge who shares your view of it somewhere to write. And because leave turns on public importance, an argument confined to settled law gives the Supreme Court no reason to take the case, while a reasoned dissent is close to an advertisement for one.

A dissent is not a win, but it is worth something. Take a typical private law case. Leave is granted in something like 5% of cases without a dissent below, and around 20% where there is one. Once leave is granted, an appellant succeeds roughly 45% of the time without a dissent, rising to roughly 55% with one.

Multiply those through. Without a dissent, the losing party's remaining position is worth about 2% of the case. With one, about 11%. The dissent is worth roughly 9% of the value of the case.

Concretely: lose at the Court of Appeal with a dissent in a $10 million case and you are not holding nothing. You are holding something worth on the order of $875,000, and your opponent's judgment is worth 89 cents on the dollar rather than 98. Both sides should be recalibrating their expectations. Many don't, because a loss at the Court of Appeal feels like the end of the road rather than a repriced position.

These figures are illustrative; they shift over time and depend on the features of the case. But the direction is stable. Our data at Lenczner Slaght, through our Data-Driven Decisions program, shows that a dissent at the Court of Appeal both materially increases the likelihood of leave and somewhat improves an appellant's odds at the Supreme Court.

Most cases will never see Ottawa. Knowing which ones might, and framing them differently from the outset, is worth far more than deciding after the Court of Appeal rules.

Part 36: When Should You Retain an Expert?

Today I want to start talking about experts. My own practice is very expert heavy, so I’ll likely have a few posts on this topic. I’ll start with timing of retaining an expert.

Most litigators typically retain an expert once they know, at least generally, what they want the expert to say. The theory of the case is settled, the pleadings are closed, and the expert's role is to (hopefully) provide independent evidentiary support for a position already taken.

That is the low-value use of an expert. In many cases, it makes sense, as there's little justification for retaining an expert in a straight-forward case that will likely settle before an expert engagement is strictly necessary. But in complex, bet-the-company cases, the high-value use comes much earlier, and it's a different purchase entirely.

In a complex case, the outcome may turn on a handful of inputs a lawyer cannot reliably assess themselves. Early in the file, a litigator's estimate of each is an informed guess dressed in confident language. And, without suggesting any bad faith, it is a guess a lawyer may be cognitively motivated to get wrong, because it supports the answer their client wants.

What an early expert provides you is not evidence. It is information: an independent estimate of a parameter you are poorly positioned to judge, at a point when the estimate can still change what you do.

That timing is key: information has value only where it can alter a decision. Once your theory is pleaded, your damages approach is set, and your client has been told the case is strong, the same opinion arrives with fewer places to go. That expert input can still alter a judge's decision, but it can't really alter your own decisions, at least to the same extent.

The corollary is that an early retainer is worth most in the scenario where it tells you something you don't want to hear. A pessimistic expert view early in the case can help you reshape the pleading, reset client expectations, and reprice settlement before costs are incurred. The same view down the line is more problematic. If you are only willing to retain early when you are confident of the answer, you are buying insurance against a risk you have already decided doesn't exist.

Again, none of which makes an early expert engagement right for every file. Experts can be expensive. Early retention front-loads that cost into the period when you are least certain the case will proceed at all. On a modest claim, the option is worth less than it costs, and the best course of action is to wait.

So when does it make sense to retain an expert early? First, where the claim is large enough to bear the early cost in expected value terms. Second, where the case has a genuine technical fork in the strategic road: that is, a technical parameter where being wrong changes the strategy. Those files are

usually identifiable early on. A case that turns on credibility rarely qualifies. A case that turns on complex technical matters or industry practice often does.

All of which is to say, think about experts early, and be prepared, in appropriate cases, to recommend engaging experts very early in the case.

Part 37: Ten Rules for Litigators using Generative AI

It feels like every week we hear a new story about AI misuse in litigation. For some lawyers, the lesson is extreme caution, which manifests as a reluctance to use generative AI at all. If you've read my content before, you won't be surprised to learn that I think that's the wrong lesson.

Generative AI is a genuinely useful technology that, used properly, can dramatically improve the quality of our work while lowering cost to clients. Refusing to use generative AI will eventually be the equivalent of refusing to use the internet or email. The problem is using generative AI poorly.

This week, I'm posting a top ten list of best practices for using generative AI in a litigation practice:

1. Learn the basics of how gen AI tools work and what they can and can't do.

2. Spend some time understanding how to prompt properly. Just as searchable legal databases made a generation of lawyers good at translating legal questions into Boolean queries, generative AI is going to turn a generation of lawyers into good prompters. Thoughtful prompting mitigates many of the risks.

3. Always check the outputs. Hallucinations are less common with advanced models, but they still happen. And even if a case isn't hallucinated, the tool might cite a real authority and misread it. Many tools link directly to sources so you can verify the basis for a statement in real time; use that.

4. Don't put privileged or confidential information into tools that aren't secure. Most free tools, and even many paid ones, train on your data. Some reserve the right to share it with third parties. Confirm neither is true before proceeding.

5. Relatedly, always consider your obligations to others. Documents produced in discovery are subject to the deemed undertaking rule. Protective orders and outside counsel guidelines may restrict AI use.

6. Know the rules where you practice. Courts take different approaches to the use and disclosure of AI-generated content in filed materials.

7. Make sure clients understand the risks of generative AI. A client's exchanges with generative AI may sometimes be relevant and producible. Some uses might attract litigation privilege, but there is little case law on that so far.

8. Preserve relevant AI exchanges. We are alive to preserving email, messages, and document repositories. Generative AI does not always make that list. It should, because retention periods on these platforms can be very short, and you don't want a spoliation problem.

9. Brief your experts. There is nothing wrong with an expert using AI for background research, proofreading, or giving expression to an opinion already formed. But an expert should not outsource the opinion itself, and should not ask questions that make it look like they did.

10. If something goes wrong with generative AI use, disclose it promptly and deal with it professionally. As in most contexts, the cover-up is worse than the crime.

Part 38: Sunk Costs and Pot Odds

Many people have heard of the sunk cost fallacy. If you've spent money and can't get it back (i.e. it's "sunk"), those sunk costs shouldn't drive your next decision. The money's gone, whatever you do. Rationally, you should only be comparing future benefits and future costs.

This fallacy comes up all the time in litigation. A party may say: "We've spent $2 million on this case. We can't walk away now." That's an understandable sentiment. But if the case has turned out to be a dead loser, the money already spent isn't a reason to keep spending. Throwing good money after bad doesn't recover the money already spent.

The sunk cost fallacy is important, but it's also an incomplete story. The reason comes from a concept familiar to poker players: pot odds.

Imagine a poker game. Suppose there's $900 in the pot and it costs you only $100 to call. You're risking $100 to win $900. If you'll win more than 10% of the time, calling is the right play, even though you'll probably lose the hand. It doesn't matter how much of that pot you put in yourself. Your earlier bets are sunk. All that matters is the price of continuing, the size of the prize, and the probability of winning it.

Litigation works the same way. Suppose you have a $1 million claim with a 30% chance of success that will cost $900,000 to take through trial. At the outset, that's a bad investment: $300,000 in expected recovery for $900,000 in spending.

Now suppose you're $800,000 in, and trial will cost another $100,000. The past spending is irrelevant. But the case is now a good investment: $300,000 in expected recovery for $100,000. Even at a 15% chance of success, you should keep going.

It's the same case with the same odds, but the rational decision is different. Not because of what you've spent, but because of what's left to spend.

Loser pays costs rules like we have in many Canadian provinces can amplify this effect. Where the loser pays, the winner recovers a share of what they've spent, so the pot grows as the case goes on. And where walking away means paying the other side's costs, folding is no longer free. That helps explain why some cases go to trial that, in hindsight, a rational plaintiff would never have started: by the end, both sides are looking at pot odds.

What does this mean for litigation strategy? A few things.

At the outset, make your best assessment of how the case is likely to unfold, what it will cost, and what course of action makes sense. But don't treat that assessment as fixed. Update it continuously as the facts, the law, and the costs evolve.

Sometimes that means cutting your losses and abandoning a case, no matter much you've already invested. But sometimes it means doubling down on a case you might never have started. Looking back from the beginning, litigating may not have been rational. Once you're deep into the case, the irrational thing can be to stop.

Part 39: The Curse of Knowledge

Every litigator knows you need to know all the ins and outs of your case. But knowing everything about a case comes with a hidden cost: the better you know it, the harder it is to see it the way the judge will.

Economists have a name for this. In a 1989 paper, Colin Camerer, George Loewenstein, and Martin Weber called it the curse of knowledge. Better-informed people can't set aside what they know when predicting how less-informed people will think, even when it's in their interest to do so.

The best-known illustration comes from a 1990 Stanford study by Elizabeth Newton. Participants were asked to tap out the rhythm of a well-known song, and a listener tried to name it. The tappers predicted listeners would get it about half the time. Across 120 attempts, listeners got just three right. The tappers heard the melody; the listeners heard only tapping.

Litigation is no different. By the time of a hearing, we may have lived with the record for years. The judge reads our factum cold. Take a typical factum in a commercial case. The team has been saying the "Amending Agreement" and the "March Notice" since the pleadings. To them, each term carries a piece of the story. To the judge, they're new labels to look up. The same thing happens with compressed chronologies, and with strong points left unsaid because they seem obvious.

Earlier in this series, I argued that effective litigators treat every case like it's going to trial, in part because preparation disciplines judgment. That's true. But the curse of knowledge is a cost of preparation, and if you're not careful, that cost only grows as preparation increases. That doesn't mean you should prepare less; it means you need to deliberately seek out an outsider's view.

So what can you do to compensate?

1. Keep the problem in mind as you prepare. Awareness alone won't cure it, but you can't address a problem you've forgotten about.

2. If your case involves complex subject matter, try to have at least one team member who isn't an expert in it. The curse of knowledge doesn't apply just to complex factual stories; it also applies to whole areas of law. Make sure your legal theory is persuasive even to someone who doesn't know every corner of that area.

3. Use a cold audience to test your argument. Lean on the newest member of the team or, if you can, someone experienced from outside your case team. They're the closest proxy you have to the bench. Give them your factum without briefing them first, and find out what confused them. Write down their first impressions: what struck them as unclear on a first read is valuable information.

So yes, you should know the ins and outs of your case. But remember: while you're humming a beautiful tune, the judge hears only the taps. The best advocates write for the listener.

Part 40: The Uses and Limits of Expected Value, Part 1

Expected value is the probability-weighted value of a course of action. Imagine you have some course of action with a number of possible outcomes. For each possible outcome, multiply its probability by its payoff, then add them up. To use a simple example, suppose a defendant faces a 40% chance of a $10 million judgment and will spend $1 million getting to trial. The expected cost of litigating is $5 million. Any settlement below

that is, in expected value terms, a good deal.

For a risk-neutral decision-maker, expected value is generally the right way to choose a course of action. Whenever I discuss a major strategic decision with a client, like settlement, I typically do some kind of expected value calculation. Because litigation is so uncertain, I find it hard to frame recommendations in any other terms.

But the expected value paradigm has its limits in practice. Or at the very least, it has to be applied carefully and contextually. Here's one example.

Imagine you act for a defendant in a piece of litigation. Holding all else equal, the bigger the claim against that defendant, the greater the expected losses from the claim. As a result, the more the defendant should rationally be willing to pay to settle it. The expected value calculation seems pretty straight-forward. But that's not how things shake out in practice.

A defendant facing a claim faces at least three different zones of risk.

Zone 1: The claim is small relative to the business. A loss would be unwelcome but absorbable. Here, expected value works exactly as advertised. The defendant is prepared to pay up to the expected cost, and not much more.

Zone 2: The loss is large enough to disrupt the business. Now it carries costs beyond the expected costs of the action itself, in the form of business disruption. Those costs tend to grow faster than the expected costs of the action itself. Business interests of the internal stakeholders usually push the same way. In this zone, it’s often rational to pay a premium over expected value for certainty.

Zone 3: The loss is catastrophic. Even a reasonable settlement on expected value would devastate the company. Here, limited liability changes the math. Owners can’t lose more than the company is worth; beyond that point, losses fall on creditors. A trial with some chance of paying nothing can beat a settlement that guarantees ruin. The company may simply have no choice but to take the risk of going to a hearing.

If you calculate expected value using only the dollar value of the judgment, it will mislead you in Zones 2 and 3. That doesn’t mean expected value doesn’t work. It means you have to account for the real costs to the business, and the limits on what it can actually lose.

Part 41: The Uses and Limits of Expected Value, Part 2

This week, I want to look at a more basic problem: how we think about probabilities. An expected value calculation is only as good as the probabilities that go into it. As I've said before, litigators should give clients quantitative probabilities rather than just descriptors like "strong case" or "reasonable chance". That's not easy, but that doesn't mean we shouldn't try.

But when you do that, one common mistake is failing to understand how probabilities compound. Behavioural economics tells us that people systematically overestimate the probability of conjunctive events: outcomes requiring several things to go right. Here’s how that can happen in litigation.

Lawyers tend to assess cases issue by issue. Suppose a plaintiff has to establish liability, causation, and damages, and counsel thinks the plaintiff has a 70% chance on each. Three issues, each more likely than not to be won. That sounds like a strong case.

In fact, it's a relatively weak case. To win, the plaintiff needs to win all three. If the probabilities are independent, the chance of winning all three is 70% × 70% × 70%, or around 34%. Add a limitations defence that the plaintiff has a 70% chance of defeating, and the chance of winning the case drops to about 24%. A case built on a series of “likely to win” issues can still be a likely to lose case overall.

The flip side matters for defendants. A defendant usually needs to win only one issue to defeat the claim. Several independent defences are worth more together than any one alone. That’s one reason building multiple paths to victory into a case theory has real value.

But there’s an important caveat. You can only multiply probabilities like that when they're independent, and often they're not. If the judge believes the plaintiff’s key witness, she may find for the plaintiff on liability and causation together. If the issues were perfectly correlated, the plaintiff’s chance would stay at 70%. The truth usually lies somewhere in between, depending on how the issues relate.

So how do you get this right? A few thoughts:

1. Break the case into the things that must all go right, and assess each separately.
2. Ask how they’re connected. Do they turn on the same witness, document, or finding? If so, they’ll tend to rise and fall together.
3. Count independent paths. Arguments resting on the same premise aren’t separate paths to victory.
4. Check your bottom-up number against your gut. If your instinct says 60% and your issue-by-issue analysis

says 30%, one of them is wrong. Figure out which.

None of that means you should throw out expected value analysis. But before using it, make sure you're thinking about the probabilities of different issues correctly.