The Hidden Cost of Falling a Step Behind AI Base Models

By Brandi Pack, Director of Innovation, UpLevel Ops

Google previewed Gemini 3 in late November, which, as I write this, is earlier this week. I spent a few minutes looking through the model scorecard, and the numbers were strong. The gains in visual reasoning were especially noticeable. At this point, though, impressive scores have become routine. Every new model lands with another set of jumps, and the pattern is starting to feel familiar.

That familiarity is what should make people pause. These frequent leaps are great for users, but they do not bode well for many of the legal tech companies trying to ride the AI wave. The ground is moving too fast. The next wave arrives before the last one settles, and the tools built on top of these models are struggling to keep pace.

This is why model chasing has become pointless. By the time anyone decides which system is in the lead, a new release shifts the baseline again. The major labs are pushing each other forward, and the direction points toward a level of intelligence that will reshape far more than product roadmaps.

That is the backdrop for this conversation. Not Gemini 3 on its own, but the momentum it represents. The AI ecosystem is accelerating, and legal teams are being asked to invest in tools that cannot evolve at the same speed as the foundation beneath them.

The Pricing Problem Nobody Wants to Admit

There is a strange disconnect playing out in the market. The base models keep improving at a pace hard to absorb, yet the prices for many AI tools are still climbing. Legal tech companies are charging more for products that add less with each release cycle. The value moves in one direction while the pricing moves in another.

In the early days, this made sense. The overlay tools were filling gaps that the base models still had. They offered guardrails, structure, and small functional boosts that made experimentation safer. Those additions gave teams early wins and earned them their place.

That era didn't last. Every time the base models level up, more of that added value gets absorbed directly into the model. The guardrails are stronger. The reasoning is sharper. The multimodal understanding is more coherent. Tasks that once required an overlay tool now run cleanly on the base model alone. With each upgrade, the overlay adds a little less and costs a little more.

It all adds up to a pattern that looks more like a bubble than a stable market. Legal tech companies want their tools to appear indispensable. The base models keep making more of those features obsolete. The pricing reflects ambition and pressure, not the reality of where innovation is happening. The mismatch is widening, and it will not widen forever.

Why Legal Tech Companies Are in a Bind

Many legal tech companies are caught in a structural bind. They took on substantial investment over the past two years, and that money came with expectations for rapid growth and a compelling story about proprietary technology. But the part of their product that delivers the most value is also the part they do not control. The base models sit elsewhere and evolve at a pace no product roadmap can keep up with.

To justify pricing, these companies need their tools to appear differentiated, even when many of the features they highlight now come standard with the newest models. Every major release absorbs more of what once set the tools apart. Roadmaps expand, promises grow, and the ground continues to move faster than the products built on top of it.

This creates a second challenge. Teams working entirely within third-party tools are effectively cut off from the new capabilities that arrive almost weekly. Features like app calling, agent mode, deep research, long context, and advanced multimodal workflows appear at the base model layer first. If the tool hasn't updated (or can't update) those capabilities, they simply don't exist for the user.

This is where the bind becomes a bottleneck. Even when these tools update their underlying model, the custom interface wrapped around it often can't keep pace. Users end up working inside a slowed-down version of the technology. Legal work exposes this immediately. The issue isn't model performance; it's the missing layer of features and tooling that the base models now provide natively. When the interface can't surface those new tools — things like advanced analysis modes, app calling embedded directly into chat, richer multimodal workflows, or deeper reasoning utilities — legal teams end up stuck with a thinner, older version of what the model can actually do. The system doesn't fail dramatically. It just quietly underperforms.

Most of the time, the problem isn't the model. It's the wrapper. In a market moving this fast, the slowest layer ends up setting the limits.

What Legal Teams Can Do Instead

There is no single playbook for a moment like this, but there are approaches that give legal teams more control and clearer visibility into what the technology can actually do.

This isn't an argument against tools. There are strong, well-designed products in the market. The issue is relying on them as the only point of access to AI.

The foundation needs to come first. Legal teams should build their initial AI capability directly on one of the leading base models. That becomes the reference point for everything else. When teams can run tasks directly on the model, they understand the true baseline. They see instantly where the model excels and where it falls short.

Only then does it make sense to evaluate tools. The process shifts from buying promises to identifying actual gaps. What is missing? What workflows need structure? What tasks require automation? Tools become purposeful, not default.

Evaluation then becomes more straightforward. How often does the tool update its model integration? How quickly does it adopt new capabilities? Which parts of the product are truly proprietary?

Pricing deserves the same clarity. If most of the intelligence comes from the base model, the premium should reflect the vendor's real contribution.

The goal isn't to avoid tools. It's to avoid being limited by them. When legal teams build their foundation at the model layer, everything else becomes easier to evaluate.

The Bubble Will Correct

The current market can feel chaotic, but the pattern is familiar. When prices rise faster than value, and the foundation shifts faster than the products built on top of it, the gap eventually closes. The correction doesn't arrive in a single moment. It shows up in buyer behavior.

Legal teams will shape this next phase more than they realize. As more organizations build directly on the base models, expectations will shift. Tools will need to offer real structural value, not cosmetic features.

The next generation of legal tech will be built by companies that embrace the speed of the models, not try to insulate their users from it. They will design for flexibility. They will assume the technology changes weekly. They will let the models do the heavy lifting and focus on the parts of legal work that truly require structure.

If they stay close to the foundation, the ecosystem will evolve around them. And that is where real transformation begins.

Learn more about the AI Advantage from **UpLevel Ops** at [UpLevelOps.com.](https://UpLevelOps.com)