AN INTERVIEW WITH...

Dave Miller, Ph.D.
Chief Scientific Officer
AustinPx

AUSTINPx™

Early formulation decisions can make or break a program, yet they're often made under significant data constraints. What are the most common mistakes sponsors make when they rush that decision, and what does it typically cost them downstream?

The most common mistake is committing to a formulation before it is understood what the molecule needs for success, usually because the clock and the material budget are both unforgiving. It shows up in a few different ways. The one people think of first is bioavailability: a team goes with a simple conventional form because it's fast and economical, and then the first PK study comes back with exposure that's too low or highly variable. But that's far from the only trap. Skimping on stability and excipient-compatibility work can let a degradation problem surface once you're already in the clinic. Ignoring properties like particle size, morphology, powder flow, hygroscopicity and manufacturability in general early can turn a formulation that looks great at the bench into something you can't actually produce at scale. There's the opposite mistake too, over-engineering a molecule onto an enabling technology it never needed, which quietly burns the time and money you didn't have to spend. And people often forget to leave themselves dose headroom, so a form that works at the starting dose can't stretch to cover dose escalation. What all of these have in common is that they tend to be discovered late, in the middle of a live program with the clinical timeline already moving. At that point the cost is rarely just the reformulation itself. It's the lost months, the extra API burn, the repeated animal or clinical work, and sometimes a mistaken read on whether the molecule is even viable. The irony is that the work to get ahead of all this up front is modest compared to what it costs to fix later, which is exactly why the early screening decision must be supported by solid, targeted datasets.

Walk us through what happens when a new molecule comes through the door- what does the intake process look like, and what information do you need from a sponsor before screening can begin?

When a new molecule shows up, our first aim is to understand it as well as the data will allow before spending any material at the bench. We start by examining all existing information from the sponsor: the structure, the solid form, whatever solubility and stability liabilities are known, the indication, and just as important, the projected clinical dose and regimen from their clinical team. We also do an EHS and handling assessment so we can bring it into the lab safely. Some sponsors show up with a full data package, but most don't, and that's expected. This is really where our preformulation capabilities shine, because we can meet the molecule wherever it is. If a sponsor arrives with little to nothing on the compound, we can run a targeted, in-depth preformulation workup and build the picture from the ground up: solid-state properties, thermal behavior, hygroscopicity, powder rheometry, and true density, alongside the core solubility and stability properties. On the bioavailability side we generate intrinsic dissolution, solubility in biorelevant media, logP, pKa, and effective permeability, plus the derived metrics like dose number, dissolution number, absorption number, and maximum absorbable dose all supplemented with a set of in silico calculations and models. Together those tell us how the molecule is likely to behave and where the real risks sit. When a sponsor already has some of this, we don't repeat work just to repeat it. We look at what exists, figure out which gaps actually stand between us and a full developability assessment, and fill those specific holes. Sometimes that's a single missing measurement, sometimes it's confirming physicochemical properties or nailing down how solubility lines up with the dose. Either way, the point of intake is to walk into screening with a clear, shared picture of the molecule and the constraints on the program, so that everything we do next is well targeted.

Thoroughness and speed are often treated as opposing forces in formulation development. How does a data-driven workflow allow you to pursue both without compromising either?

We don't treat thoroughness and speed as opposing forces because the workflow is built to answer the questions that actually drive the formulation decision, and to answer them in the right order. A data-driven development assessment lets us front-load the small number of measurements and models that determine which path a molecule should take, rather than running an exhaustive battery on every candidate. Once the critical questions are resolved, effort flows to where the data say it is needed. That focus is what allows comprehensiveness and speed to coexist: we are thorough about the decisions that matter and efficient about everything else. API-sparing designs, predictive modeling, and parallel screening compress the timeline further. The

result is not a lighter version of a comprehensive program, but a smarter one that reaches a defensible answer with less material and less time.

There's a "pivotal branch point" in development assessment - the fork between a speed-to-clinic candidate and one that needs bioavailability enhancement. How do you know when you've gathered enough data to make the call confidently?

The fork between a speed-to-clinic candidate and one that needs bioavailability enhancement turns largely on the relationship between biopharmaceutical properties and the projected clinical dose. We reach that call confidently when the biopharmaceutical picture is internally consistent: biorelevant solubility, intrinsic dissolution, logP, permeability, and the derived metrics such as dose number and maximum absorbable dose all point in the same direction, and modeling corroborates the measured data. When they agree, the decision is clear and we move. When they conflict, that disagreement tells us precisely which additional experiment will resolve it, so we run that one rather than continuing to gather data broadly. Confidence does not come from having every possible measurement; it comes from convergence among the measurements that matter. Knowing which those are, and recognizing when they have converged, is where experience is critical.

> When a new molecule shows up, our first aim is to understand it as well as the data will allow before spending any material at the bench.
>
> Dave Miller, Ph.D.

API scarcity is a constant pressure in early development. How has your approach to study design evolved to extract more insight from less material, and where do you see the biggest opportunity for further efficiency?

API scarcity has shaped everything about how we design early studies. The guiding principle is to extract the maximum decision-relevant information from the minimum amount of material, which pushes us toward miniaturized and micro-scale methods, experimental designs that avoid redundant work, and use of predictive modeling to decide which experiments are worth the material spend. In silico assessment lets us rule approaches in or out before consuming API, so the material we do spend goes to the questions that genuinely require an experimental answer. Analytical methods that need only small sample masses stretch a limited supply further. The biggest opportunity ahead, in my view, is continuing to tighten the modeling-to-experiment loop, so each round of data sharpens the predictions and further reduces the material needed for the next decision. Doing more with less is not a constraint we tolerate; it is a discipline the workflow is built around.

Amorphous solid dispersions are notoriously difficult to screen because no single platform fits every molecule. How does having spray-drying, HME, and KinetiSol in-house change the quality of the decisions you're able to make?

Amorphous solid dispersions can be challenging to screen precisely because no single manufacturing platform is right for every molecule. A compound that degrades at the temperatures required for hot-melt extrusion may be an excellent spray-drying candidate; a molecule that resists both may still succeed with KinetiSol, which reaches the amorphous state through very short thermal processing times and without solvent. Having spray drying, hot-melt extrusion, and KinetiSol together under one roof changes the quality of our decisions because we are not trying to force a molecule onto one platform. We let the science decide which technology delivers the best performance, stability, and manufacturability, and we can screen across them in parallel. AustinPx offers the most comprehensive ASD screening capability of any CDMO in the world, and for a poorly soluble molecule that breadth is often the difference between a viable enabled formulation and a dead end.

The AustinPx team has experience on the sponsor side. How does that perspective- having lived the timeline and budget pressure firsthand- actually change the way you run a program?

The company spent years on the sponsor side, so we have lived the pressures our clients are under rather than just hearing about them. We have faced the same limited API, the same board-driven timelines, and the same need to make a defensible decision without the luxury of unlimited data. That experience changes how we run a program in concrete ways. It makes us disciplined about spending material and time only where they change a decision. It makes us communicate in the terms a sponsor's team actually has to answer to, tying the science back to the program's clinical and financial realities. And it gives us the judgment to know when the data are good enough to act and when they are not, which is exactly the call that early development lives or dies on. High science and high service, informed by having been in the sponsor's seat.