Signal from Noise

July 25, 2026

AI AT PHASE VALUE – TAKING ON CLINICAL DRUG TRIALS

Kent Fung

VICE PRESIDENT, MARKET INTELLIGENCE

The previous *Signal From Noise* discussed the process by which potential new medicines and treatments are identified. Yet that only describes a third of the path. Before a drug developer – and patients – can benefit from a new drug, the company must first determine whether it works, and, just as importantly, prove it to regulators and to the medical community.

Although the clinical-trial and the regulatory-approval processes might seem less dramatic or glamorous than drug discovery, they are arguably just as challenging and complicated. They certainly are time-consuming and costly. Traditionally, it can take six to eight years and sometimes billions of dollars to shepherd a drug candidate through the three main phases of the clinical trial process, and another year or two to convince regulators to approve a drug candidate for the market.

Here once again, pharmaceutical companies are hoping that AI can not only shorten and streamline the process, but improve the odds of success while also lowering costs.

While we briefly touched upon the major phases of clinical drug trials in our last issue, it is worth going over in more detail now to provide industry outsiders with a better understanding of what makes clinical trials such uncertain, complex, and challenging endeavors.

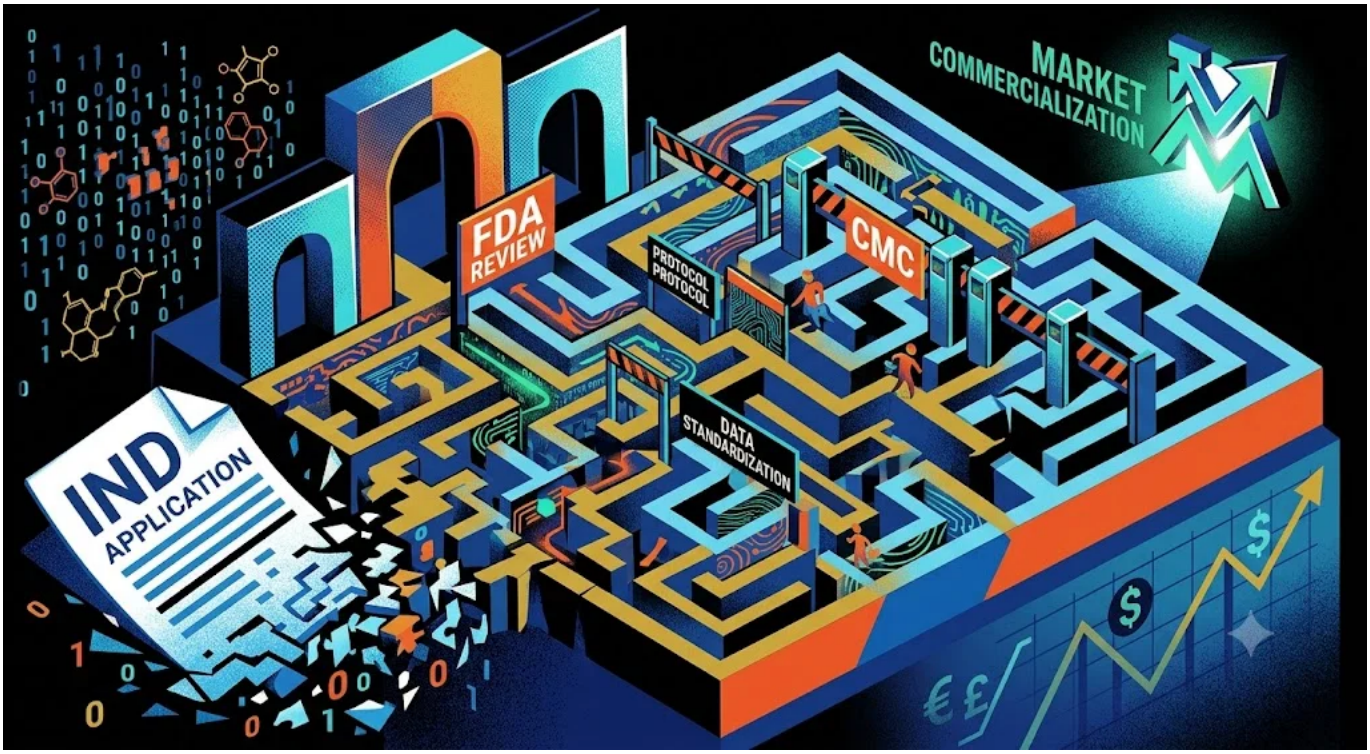

Source: Fundstrat/Gemini

IND – Investigational New Drug status

Before any company can begin testing a drug candidate on humans, it must get regulatory approval. And while we will focus on the U.S. process, European and Japanese regulatory standards are roughly as rigorous and involved.

Human trials cannot begin until the potential treatment attains Investigational New Drug status. This begins with submitting documentation to show regulators that the proposed medicine is likely safe for humans, often by showing results from animal testing or lab assays.

Companies must also demonstrate adequate chemistry, manufacturing, and controls, or CMC. CMC involves proving that a company can manufacture the drug to be tested at a sufficiently high quality. This includes showing that the medicine can be made with sufficiently pure, sterile quality, and in a form that is stable. Most importantly, the company must demonstrate that it can consistently achieve such manufacturing quality time and time again.

IND approval is also contingent on convincing regulators that the scientists involved have a sound plan on how to run the clinical trials. Many of the decisions about the clinical trials of a given drug candidate must be made well in advance, when submitting an application for regulatory approval to begin human trials.

Phase I

The Hippocratic oath requires aspiring doctors to swear that first they will do no harm. The initial stage of human clinical trials involves establishing not that a potential treatment works, but that it does no harm – or at least that any harm is either minimal or highly unlikely. This involves testing a drug candidate on a small (20 to 100) number of healthy volunteers, seeking to establish safety, toxicity, and maximum tolerated dosage, while acquiring data on what the human body does with a drug – how it absorbs, metabolizes, and excretes it. (*Note: This field of study is known as pharmacokinetics, or PK. A closely related field, Pharmacodynamics or PD, involves studying how a drug changes or affects the body*.)

Phase I trials exist because a substance that seems safe when tested on animals can obviously turn out to have a very different effect on humans. Even if the drug turns out to be essentially non-toxic, Phase I is also when a company might discover that the human body cannot effectively absorb the medication, instead excreting it before it can have any meaningful effect.

A notable exception is made for Phase I trials of some proposed cancer treatments. Often, cancer treatments like chemotherapy involve administering medicines that are known to be toxic. (Chemotherapy is essentially a bet that, even though the treatment is toxic to the patient, it will hopefully be even more toxic to the cancerous cells.) Because such drug candidates cannot ethically be tested on healthy volunteers, they are instead tested on volunteer patients in the late stages of cancer who have already exhausted other treatment options.)

Phase II

The IND and Phase I stages are essential, but for a drug developer, Phase II is where researchers find out whether they've wasted their time and effort, because this is when the trials seek to discover whether the drug actually works on the targeted

disease or ailment. The group of test subjects here becomes slightly larger – 100 to 300 patients, all of whom suffer from the target condition.

Testing a drug against a disease is not as simple as that, however. Not all patients with a disease are the same. For example, some treatments that were thought to be broadly effective were later found to be significantly less effective on women because a disproportionate percentage of test subjects were men. A proposed drug for, say, Alzheimer's, might work on patients in the early stages of the disease but have little effect on patients who have already progressed to advanced dementia. Others are effective on Alzheimer's patients who test positive for certain toxic proteins but less so on those who lack those specific biomarkers.

The general objective – both for health and investment purposes – is to develop a treatment that works on most if not all patients with a given illness. Achieving that objective begins with Phase II testing. Assembling a pool of test patients that is sufficiently diverse in meaningful ways can help companies determine whether that objective is likely to be reached before going through the far more costly and time-consuming Phase III.

Phase II is also when researchers try to pinpoint the best way to administer a medication: one strong dose a day or an even stronger one once a week? Many smaller doses interspersed throughout the day or week? The regimen must be strong or intense enough to have an impact, but too intense a dosage regiment could make the side effects so severe as to deter patients from continuing.

Further complicating the trial process is the human element. Beginning with Phase II, clinical trials often require many, many tests — more than would be necessary for a patient taking an already-approved treatment. Instructions for dosage, if self-administered, are often precise, requiring that patients learn – and adhere to – a sometimes complicated process, perhaps in addition to whatever existing treatment regimen they had already undertaken before the trials.

Then there's the placebo effect. Although a double-blind test involving placebos is the gold standard when testing for efficacy, when the ailment being treated is serious or

life-threatening, recruiting subjects to test an experimental treatment becomes far more challenging if they learn that they might be given a placebo.

It's also worth noting that generally, patients are already taking an existing treatment for their condition. While from a scientific standpoint, the clearest evidence of success would be achieved by asking patients to stop taking that treatment, this is ethically unfeasible, so it can be difficult to determine how much improvement is due to the drug candidate, how much is due to the existing treatment test patients are still taking – and perhaps whether there is a synergistic effect.

Phase II is rarely a long-term trial stage. Because success in treating a condition is often a matter of years rather than weeks, researchers often decide success or failure at this stage by using surrogate metrics. For instance, if a drug candidate is intended to reduce heart-disease mortality, Phase II might instead assess whether the candidate lowered blood pressure or cholesterol levels. That brings us to …

Phase III

One of the reasons Phase III tends to be such a big deal is that it tends to take much longer – often running into multiple years. This enables pharma companies and regulators to rely less on the surrogate datapoints mentioned above. Yet this is just one of the reasons why Phase III trials are so challenging and expensive. Convincing patients to stay in trial and submit to frequent testing and adhere to the strict regimen for such a prolonged period can be difficult. Historically, companies have resorted to making the sample pool larger than strictly necessary even if this adds to the expense, in the expectation that some patients will drop out or start to deviate from the established protocols.

That's not the main reason why Phase III significantly expands the pool of test patients. As anyone who's ever taken Statistics 101 knows, the larger the sample group, the more reliable the conclusions are likely to be. In this third phase, the sample size gets considerably larger – anywhere from 1,000 to 3,000 patients, though in some cases, Phase III trials have involved as many as 10,000 patients.

Determining how many patients to test becomes a difficult decision. More patients means increased logistical complexity and significantly higher costs, but particularly if

a drug candidate is expected to be only slightly more effective than an existing treatment, using too small a sample size might not yield results conclusive enough to pass regulatory muster. Large sample sizes are also the key to determining whether a candidate causes rare – but serious or even fatal – side effects. If a proposed treatment causes a serious side effect in 1 out of 1,000 people, you are unlikely to discover this with a sample size of 100.

(You know those seemingly bizarre and alarming disclaimers one hears at the end of prescription-drug commercials on television? They're frequently there because of discoveries made during Phase III.)

Finding a sufficiently large pool of patients in a single city is often impossible, so Phase III trials are often simultaneously conducted in multiple locations. This geographical diversity also can help increase the reliability of the results. Why? Because it helps to ensure that the treatment is tested across a range of genetic profiles, diets, ethnicities, and regional standards of care. This also reduces the likelihood of some localized environmental factor – idiosyncratic pollution, for example – skewing the findings.

Phase II trials are often called the "valley of death" because the failure rate at that stage is the highest. Yet Phase III is commonly viewed as the most challenging – "pivotal," to use an industry term – because it is the part that regulators rely on the most to decide on approval.

What companies are using AI to help?

Certara $CERT. Certara's offerings focus on biosimulation, using predictive AI trained on toxicology databases and on pharmacokinetic and pharmacodynamics datasets to predict toxicity and other potential hazards. In the IND part of the process, this helps drug developers determine initial test doses that are likely to be safe, yet high/non-conservative enough to be meaningful, reducing or in some cases even eliminating the need for costly and time-consuming animal testing. This functionality also yields benefits in Phase I, which help to define dose-escalation protocols and predict drug interactions; in Phase II, when it can assist in predicting how different

sub-populations will respond to various doses and dose protocols (frequency, timing, etc.)

Perhaps more importantly after trials actually begin are Certara's Simcyp physiologically-based pharmacokinetic (PBPK) simulator and its Phoenix PK/PD platform. These enable Certara to help drug developers with virtual-population modeling, which enable drug companies to simulate responses across a range of patient demographics and predict likely drug interactions. In certain circumstances, regulators will accept the results of these simulations as evidence to support label claims without physical trials.

The company announced a partnership with Nvidia on July 7 that would incorporate the latter's BioNeMo Agent Toolkit – specifically, its specialized AI agents – to better assess and analyze scientific models and datasets. Certara's offerings also include generative AI and AI-assisted clinical data standardization, both which make it easier to prepare regulatory submissions. The company claims more than 2,600 clients in the biotech/pharmaceutical industry, including major companies like Pfizer and Novartis.

Veeva $VEEV. Veeva's workflow-software offering proposes the leveraging of AI agentic capabilities to help with many of the regulatory and administrative requirements involved with clinical trials. (Many of the capabilities are in the rollout phase as of this writing but not yet released.) In the IND phase, this includes assisting with drafting and organizing the many documents, datasets, and safety profiles that must be submitted in order to get approval for human trials, including, but not limited to, manufacturing documentation and proposed trial protocols. Veeva's Vault RIM also assists in registration tracking and regulatory correspondence. Later on, Veeva software is used to monitor protocol compliance, validate results, and flag imminent potential risks with regard to data quality and safety. Veeva also offers an AI-powered document management system that helps integrate trial records with hospital Electronic Health Records.

Veeva also operates an AI-powered integrated pharmacovigilance safety platform. This platform tracks reported complaints about side effects from both drug candidates and approved drugs all around the world, flags noteworthy patterns, and

automatically handles mandatory regulatory reporting when such patterns emerge. Veeva's customers include Merck, Roche, Bayer, Bristol Myers Squibb, Teva, and BioMarin.

IQVIA $IQV. IQVIA is not just an AI-driven SaaS company. In fact, the bulk of its revenues come from its business as a contract research organization (CRO) – in other words, part of a long-established industry that exists to run outsourced clinical trials on behalf of drug companies. (The company is in the early stages of starting a contract AI drug-discovery business as well.) IQVIA uses AI to streamline the complex and multipronged regulatory/administrative/bureaucratic processes required before beginning clinical trials. Natural language processing and generative AI agents help draft documents and standardize report formats. IQVA also helps to identify possible trial design flaws and future compliance red flags.

Once clinical trials have begun, IQVIA offers an Orchestrated Clinical Trials platform now powered by AI agents to help optimize test-subject recruitment and enrollment, identify the best sites to hold trials, identify what types of diversity goals are needed and how to achieve them, and fine-tune inclusion and exclusion criteria. The company's AI agents can also help clean up data formats, detect protocol deviations and data anomalies in real time, and thus shorten the timelines. IQVIA's capabilities stem in large part from a core database of 1.2 billion anonymized patient records – roughly 56 petabytes (and counting) of longitudinal healthcare data. As with others, IQVIA has collaborated with Nvidia and Amazon Web Services to develop these capabilities. Its customers include most of the largest global pharmaceutical companies by revenue, including the likes of Pfizer, AstraZeneca, and Eli Lilly. Notably, IQVIA in August 2025 formed a partnership with Veeva, which means that the two companies use each other's AI capabilities.

Tempus AI $TEM. Tempus is an oncology-focused company that combines real-world data with information from tumor-tissue repositories (cells cultured from tumor samples) and a database of 38 million research records and 7 billion clinical notes. These inform an AI model for drug discovery and also help to identify the likely therapeutic sensitivity of proposed cancer treatments.

Tempus's focus on and close work with oncology practices enables the company to apply AI to quickly identify and pre-screen potential patients suitable for various clinical trials. Cincinnati's TriHealth health system credited Tempus for trial enrollment improvements of 50% or more, while Sermonix Pharmaceuticals asserted that Tempus helped to shorten trial enrollment times by as much as 10 months. The company also uses this asset to speed and improve trial site selection.

As clinical trials proceed, Tempus leverages AI to deliver real-time recommendations on choosing and identifying biomarker-defined patient subtypes and cohort selection. More importantly for a company focused on cancer research, Tempus combines AI with its specialized cancer dataset to build synthetic control groups, which as we explained above, are essential for cancer-drug research that ethically cannot use placebo groups. According to its client AstraZeneca, Tempus's AI technology has helped boost the likelihood of technical success by an average 5 percentage points. Among the customers or partners in the pharmaceutical industry using Tempus technology, Bristol Myers Squibb, AstraZeneca, and GSK stand out.

The companies described thus far have been more pure-play in the new specialty of AI-driven clinical trial products and services. Yet two larger conglomerates – with other, more established businesses and sources of revenue – are active in this space as well.

Medidata is a subsidiary of French software conglomerate Dassault Systèmes $DASTY. (Dassault also makes design and simulation software for the aerospace, defense, and automotive industries.) With Medidata, the company uses its AI Rave platform to study a broad range of metrics across 38,000 clinical trials and 12 million patients across 2,300 customers, helping clients pre-test potential trial protocols and identify areas that might cause operational bottlenecks further down the road. This also assists with Medidata's ability to create synthetic control groups that it claims can reduce – or in some cases, eliminate – the need for placebo control groups in Phase II trials.

During trials, Medidata's data-management capabilities include smart integration of electronic health-record hospital data. Importantly, it can also improve the patient experience and trial reliability by pairing AI-driven analysis with real-time

sensor/wearables data. This enables decentralized trials, facilitates patient monitoring, and makes test-subject retention more likely.

The company's expertise in data management comes into play to quickly identify data anomalies in real time and to ensure that data from multiple trial sites is kept in a uniform, standardized format for reporting and regulatory-approval applications. Medidata clients include Sanofi, the CROs Worldwide Clinical Trials and Caidya, and the National Cancer Institute

Oracle $ORCL Life Sciences. Oracle's clinical-trial offering includes agentic AI applied to longitudinal records from over 129 million anonymized patients within the Oracle Health network – its electronic health record platform used by an estimated 1,200 U.S. hospitals (about 20% of the U.S. market). These capabilities are used to manage workflows during the IND application stage and in all three phases of the clinical trial process, helping to optimize and facilitate protocol decisions, identify potential operational bottlenecks, and format data for regulatory submissions.

Oracle Health Sciences also uses AI to construct synthetic test-subject data (i.e., simulated placebo groups) to support control testing, and predict site recruitment velocity in Phase III trials. Also notable is Oracle's Argus pharmacovigilance engine, an industry standard used to track millions of patient complaints annually from around the world about side effects from approved drugs. This is both powered by and helps to inform the company's AI capabilities.

Oracle sells its AI offerings as being able to view customer data, third-party sources, and 129M+ de-identified longitudinal records as a unified whole. The platform thus purportedly enables pharmaceutical companies and CROs to apply Oracle's agentic reasoning holistically throughout their workflows, in the trial and approval process – and beyond.

Conclusion

Part I of our deep dive into the use of AI in pharmaceutical research and development highlighted the "gee whiz" aspect of the AI revolution, in which scientists and doctors seek to make the kinds of breakthroughs that garner headlines.

Part II covered the part that gets a lot less of the glory, even though we would argue that it is at least as important, if not more so – not the least because it's such a costly yet uncertain part of the process. After all, as a fictional attorney once said while preparing for a different kind of trial in a well-known courtroom drama, "It doesn't matter what I believe, it only matters what I can prove." The hope is that AI can make that second part easier, faster, and less expensive for drug developers.

It could be argued that both parts of our deep dive discuss whether AI is a new aspect of the picks and shovels play for pharmaceutical companies. Unlike the use of AI in drug discovery, applying AI to the clinical trials and regulatory approval processes appears to be providing confirmed benefits, particularly in terms of speeding up the trials process and making test-subject recruitment more efficient. To us, the full benefits have yet to be fully or conclusively quantified, but we have nevertheless identified the companies we believe to be doing the most in this area. As always, *Signal From Noise* should not be used as a source of investment recommendations but rather ideas for further investigation. We encourage you to explore our full *Signal From Noise* library, which includes a fresh look at the Magnificent Seven, deep dives on why investors should care about Gen Xers, the recent memory chip gold rush, the future of malls, and drone warfare. You can also find our take on space-exploration investments, defense stocks, and the business of farming.

Disclosures

This research is for the clients of Fundstrat Global Advisors only. For additional information, please contact your sales representative or Fundstrat Global Advisors at 150 East 52nd Street, New York, NY, 10022 USA.

Conflicts of Interest

This research contains the views, opinions and recommendations of Fundstrat. At the time of publication of this report, Fundstrat does not know of, or have reason to know of any material conflicts of interest.

General Disclosures

Fundstrat Global Advisors is an independent research company and is not a registered investment advisor and is not acting as a broker-dealer under any federal or state securities laws.

Fundstrat Global Advisors is a member of IRC Securities' Research Prime Services Platform. IRC Securities is a FINRA registered broker-dealer that is focused on supporting the independent research industry. Certain personnel of Fundstrat (i.e., Research Analysts) are registered representatives of IRC Securities, a FINRA member firm registered as a broker-dealer with the Securities and Exchange Commission and certain state securities regulators. As registered representatives and independent contractors of IRC Securities, such personnel may receive commissions paid to or shared with IRC Securities for transactions placed by Fundstrat clients directly with IRC Securities or with securities firms that may share commissions with IRC Securities in accordance with applicable SEC and FINRA requirements. IRC Securities does not distribute the research of Fundstrat, which is available to select institutional clients that have engaged Fundstrat.

As registered representatives of IRC Securities, our analysts must follow IRC Securities' Written Supervisory Procedures. Notable compliance policies include (1) prohibition of insider trading or the facilitation thereof, (2) maintaining client confidentiality, (3) archival of electronic communications, and (4) appropriate use of electronic communications, amongst other compliance related policies.

Fundstrat does not have the same conflicts that traditional sell-side research organizations have because Fundstrat (1) does not conduct any investment banking activities, (2) does not manage any investment funds, and (3) our clients are only institutional investors.

This research is for the clients of Fundstrat Global Advisors only. Additional information is available upon request. Information has been obtained from sources believed to be reliable, but Fundstrat Global Advisors does not warrant its completeness or accuracy except with respect to any disclosures relative to Fundstrat and the analyst's involvement (if any) with any of the subject companies of the research. All pricing is as of the market close for the securities discussed, unless otherwise stated. Opinions and estimates constitute our judgment as of the date of this material and are subject to change without notice. Past performance is not indicative of future results. This material is not intended as an offer or solicitation for the purchase or sale of any financial instrument. The opinions and recommendations herein do not take into account individual client

circumstances, risk tolerance, objectives, or needs and are not intended as recommendations of particular securities, financial instruments or strategies.

The recipient of this report must make its own independent decision regarding any securities or financial instruments mentioned herein.

Except in circumstances where Fundstrat expressly agrees otherwise in writing, Fundstrat is not acting as a municipal advisor and the opinions or views contained herein are not intended to be, and do not constitute, advice, including within the meaning of Section 15B of the Securities Exchange Act of 1934. All research reports are disseminated and available to all clients simultaneously through electronic publication to our internal client website, fundstrat.com. Not all research content is redistributed to our clients or made available to third-party aggregators or the media. Please contact your sales representative if you would like to receive any of our research publications.