# A beautiful system

*-- how science becomes medicine --*

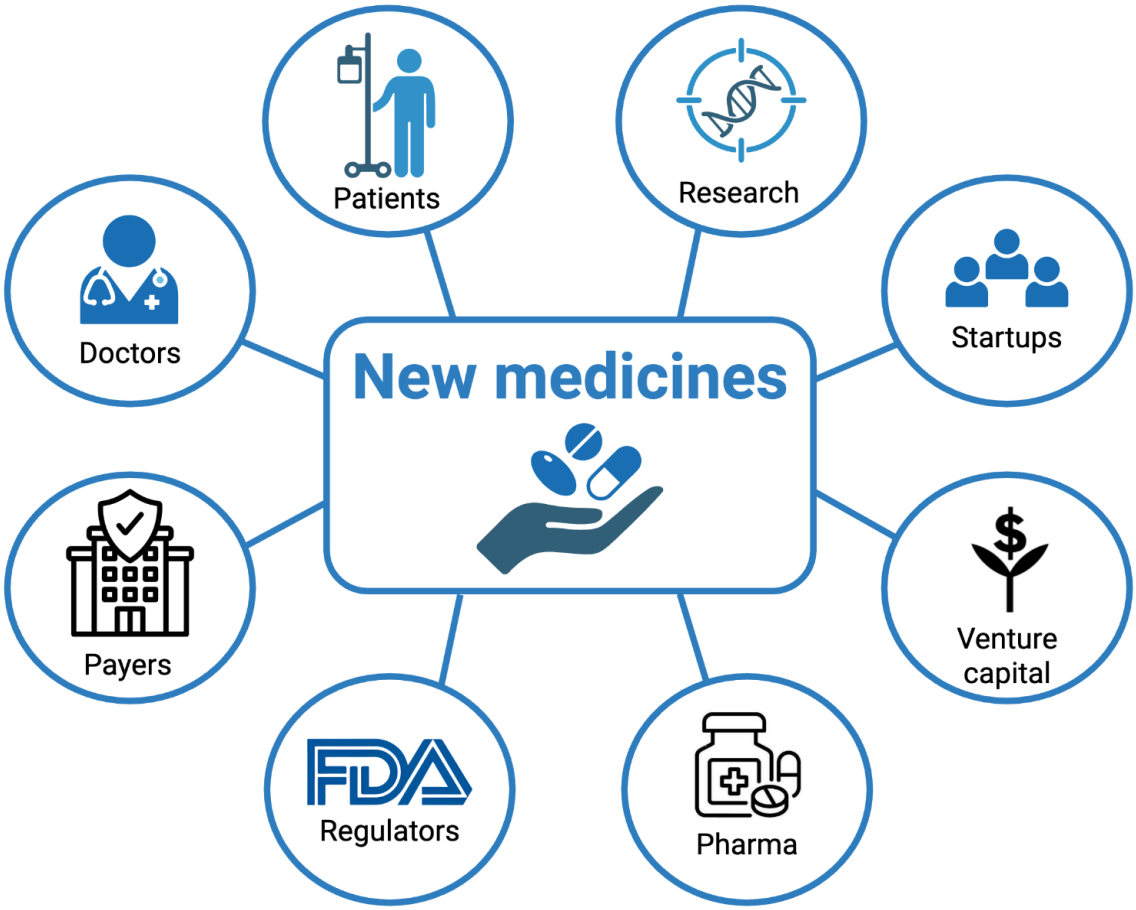


***Javier Romanelli -- last updated in April 2026***

# Introduction

Most people have no idea how a drug gets made. They watch drug ads on TV and fill their prescriptions, but the system connecting academic scientists to that prescription bottle is a blurry picture at best. Recently, a smart friend of mine who runs a tech startup plainly asked me -- how does a drug *actually* get made? He understood startups, venture capital, product-market fit, and acquisition dynamics. But the biotech version of all that was completely foreign to him. How do scientific discoveries become drugs? How do you patent a molecule? Why do clinical trials cost hundreds of millions of dollars? Who decides whether your insurance covers a new treatment? And who are all the players in between?

First things first: what is a drug? At its most fundamental level, it's a molecule that enters the body, finds a specific biological target (usually a protein), binds to it, and produces a desired effect. That's it. The target might be an enzyme, a receptor on a cell surface, a signaling molecule floating in the blood, or a stretch of DNA. The molecule might be small (like aspirin, which inhibits cyclooxygenase enzymes) or large (like Humira, a monoclonal antibody that binds TNF-alpha). But the core logic is usually the same: molecule meets target, target's function changes, downstream biology shifts, patient feels better.

Early stages of drug discovery are about finding the right molecule for the right target. Clinical trials are about determining if a molecule hits its target in living humans and produces the predicted effect, all without causing unacceptable harm. Surprisingly, many approved drugs (that is to say, they work in patients) don't have their molecular mechanisms precisely mapped out. This is especially true for older drugs, because the field was very different a few decades ago and clinical development cycles span roughly a decade.

Understanding the biopharmaceutical industry is important because it generates over $500B in annual US revenue, employs millions of people globally, and directly determines the quality and length of human life. It's one of the most capital-intensive, heavily regulated, and scientifically complex industries on the planet. Perhaps more simply (and personally), it's important for people to be properly informed to decide why they should (or shouldn't) trust a doctor's prescription. Yet there's a gap between what the industry produces and what people understand about how it works, even for many who work in the industry.

You might reasonably ask why anyone would read a long writeup in 2026, when you can ask an AI to explain any of it in seconds. The answer is that information and understanding are different things. In my opinion, there's something about following one story from beginning to end, watching each stakeholder show up in context (rather than in isolation), that a chatbot format doesn't replicate. I also wrote this as a personal exercise to spell out my mental model of the biotech ecosystem in one living document I can revisit and refine as the world and my thinking evolve.

This guide is my attempt to map the entire system as compactly and clearly as I can. I want to cover every major stakeholder: scientists in academic labs, biotech startups, venture capitalists, big pharma, regulators, doctors, patients, and the payers who sit between all of them. I want to show how their incentives align (and occasionally don't) and explain the economics, the intellectual property, the clinical development process, and the commercial dynamics. And I want to do all of this through a concrete example, so you can see each player show up in a real story rather than as an abstraction.

The example I chose is **Humira** (adalimumab), the best-selling drug in pharmaceutical history with over $200B in cumulative revenue. Its journey starts in a Cambridge University lab in the early 1990s, passes through a German chemical conglomerate, gets acquired by an American healthcare company, spins off into its own pharma giant, fights a patent war, transforms the lives of millions of patients with autoimmune diseases, and eventually faces competition from biosimilars. Every stakeholder in the ecosystem shows up in this story, making it the right vehicle for this guide.

At the end, I'll briefly cover two other drugs worth knowing: the GLP-1 drugs like Ozempic that are reshaping the obesity treatment landscape, and Keytruda, the cancer immunotherapy that almost got shelved before becoming the world's current best-selling drug.

In my opinion, the biotech ecosystem is one of the most elegant systems humans have built. Not because it's simple (it's not) and not because it's perfect (it's not), but because it aligns a remarkably diverse set of incentives toward a single outcome: getting effective treatments to patients. Scientists want to discover. Entrepreneurs want to build. Investors want returns. Doctors want options. Patients want to feel better. Regulators want safety. And somehow, imperfectly but reliably, the system channels all of that into medicines that work. It's a genuinely positive-sum game -- when a drug succeeds, everyone along the chain benefits.

For reasons I'll highlight throughout the guide, I find this machine genuinely beautiful (and I hope you will too). This guide is written so that any curious reader (from high school student to biotech venture capitalist) can find something valuable in it. *Let's take a look at how the machine works.*

# Part 1 -- the birth of drugs

## Where are drugs invented?

Most drugs don't originate in pharmaceutical companies. So, where do they come from? A 2023 analysis found that ~65% of leading FDA-approved drugs originated from external sources (biotech companies, academic labs, or collaborations between the two), while only about 28% were discovered and developed entirely in-house by large pharmaceutical companies. In oncology, the picture is even more skewed: large pharma originates under 20% of first-in-class cancer drugs, while academia and small biotech companies originate over 40% each. The pattern is consistent: the intellectual raw material comes from publicly funded research, and the private sector's role is to translate it into approved therapies.

Why? Because discovery and development require fundamentally different capabilities. Academic labs operate on curiosity, government grants (NIH, the largest funder of biomedical research in the world, distributes $48B annually in the US as of 2026), and the incentive to publish papers for further funding. They're optimized for asking new questions and finding new biology -- not for turning those findings into drugs. The latter requires regulatory expertise, clinical trial infrastructure, manufacturing capabilities, and lots of capital. The ecosystem solves this mismatch by creating a handoff mechanism: academic discoveries flow into biotech startups (through licensing deals and spinout companies), which then either develop the drug themselves or partner with pharmaceutical companies.

Public investments flow to thousands of universities, medical schools, and research institutes as competitive grants. These fund the salaries of graduate students and postdocs, the cost of materials and equipment, and the operational expenses of running a lab. In return, grantees are expected to publish their findings in peer-reviewed journals, making the knowledge publicly available. This is the foundation on which everything else is built. Without NIH-funded basic research, we wouldn't have a pipeline of druggable targets, validated disease mechanisms, and novel therapeutic modalities. A recent study found that 354 of 356 drugs approved by the FDA between 2010 and 2019 were supported by NIH-funded research, either directly (through grants that contributed to the drug's development) or indirectly (through grants that established the foundational science of an eventual drug). Both the ecosystem's dependence on public investment in basic research and the return on investment are enormous.

The technology transfer office (TTO) at a university is the institutional bridge for the handoff. When a scientist makes a discovery with commercial potential, the TTO files a patent on behalf of the university, then either licenses the IP to an existing company or a spinout by the inventor(s). In 2020 alone, this handoff mechanism led to the creation of over 1,000 startups nationwide.

Though the exact spinout terms vary by university, they typically look something like this: the university retains ownership of the patent and grants the startup an exclusive license; in exchange, the startup gives the university a royalty on future revenues (typically <5%), an equity stake (typically a few %), and possibly milestone payments tied to clinical and regulatory achievements. Stanford and MIT tend to be more founder-friendly, while some university systems have more rigid, less negotiable structures. These economics matter because they determine how much of the value created by academic discoveries flows back to the institutions that funded them, and how much remains with the company taking on the commercial risk.

## The birth of Humira

Humira's story begins in exactly this academic-to-commercial pattern. Following the invention of phage display by George P. Smith in 1985, Gregory Winter at the Medical Research Council's (MRC) Laboratory of Molecular Biology in Cambridge, England, transformed it into a powerful platform for selecting fully human therapeutic antibodies. The idea was to use bacteriophages (viruses that infect bacteria) as a screening platform for antibodies.

To understand why this matters, it's useful to understand how antibody drugs work at a molecular level. An antibody is a Y-shaped protein made by your immune system. It works by binding to a specific target (called an antigen) with extraordinary precision. The tip of each "arm" of the Y contains a variable region that determines

what the antibody sticks to, and these variable regions are what scientists engineer to target specific disease-causing proteins. The rest of the antibody structure handles the downstream immune response: recruiting other immune cells, triggering complement activation, or simply blocking the target from functioning. "Monoclonal antibody" (mAb) means a population of identical antibodies, all targeting the same thing, produced from a single cell line. This is what makes them so powerful as drugs: you can design a molecule that specifically targets one protein in the body and blocks its activity, without (in theory) touching anything else.

Before phage display, therapeutic antibodies were typically made by immunizing mice with a human disease target, harvesting the mouse antibodies that bound the target, and then "humanizing" them (replacing the mouse protein sequences with human ones so the patient's immune system wouldn't reject them). This worked, but it was imprecise and the resulting antibodies were chimeric (part mouse, part human). Infliximab (Remicade), an early TNF-alpha inhibitor, was made this way. Phage display allowed scientists to bypass the mouse entirely and select fully human antibodies from the start. The technique is elegant: you create a library of billions of phage particles, each displaying a different human antibody fragment, then you pan them against your target protein. The phages that bind get amplified. You repeat this selection process several rounds, and what emerges is a highly specific, fully human antibody that was never anywhere near a mouse.

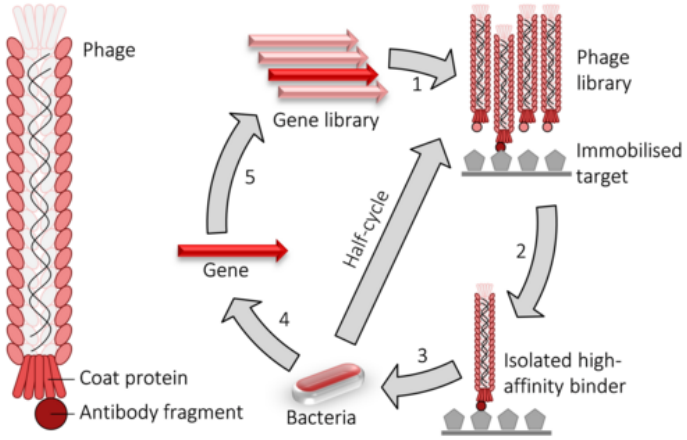


In 1989, Winter co-founded Cambridge Antibody Technology (CAT), an MRC spinout company, to commercialize phage display for therapeutic antibody development. This is the first handoff -- from a government-funded academic lab to a small biotech company. Winter's work was foundational enough that he'd later win the 2018 Nobel Prize in Chemistry (jointly with George P. Smith) for the phage display of peptides and antibodies.

CAT was a small company taking a platform technology and figuring out how to turn it into drugs. It couldn't (yet) manufacture drugs at scale, run large clinical trials, or sell drugs. The company's entire value proposition was the phage display library that Winter had built -- the ability to fish the right antibody out of a pool of billions. That capability, with the right partnership, turned out to be *enormously* valuable.

In 1993, CAT entered a research collaboration with BASF Bioresearch Corporation, the pharmaceutical research arm of the German chemical giant BASF. The target was tumor necrosis factor alpha (TNF-alpha), a cytokine (immune signaling molecule) that was emerging as a central driver of inflammatory and autoimmune diseases. Brief note on the biology: TNF-alpha is produced by immune cells and normally helps the body fight infections. But in autoimmune diseases like rheumatoid arthritis, the immune system is misdirected and attacks the body's own tissues. TNF-alpha drives this attack, causing chronic inflammation that destroys joints, damages skin, and inflames the gut. The therapeutic hypothesis was simple -- if you can block TNF-alpha, you can reduce the inflammation and stop the tissue damage. This had already been validated by earlier drugs like infliximab. The question was whether a fully human antibody could do it better.

Using Winter's phage display technology, CAT scientists identified a candidate antibody designated D2E7 that bound TNF-alpha with high affinity and specificity. It was the first fully human antibody to target TNF-alpha -- in theory, this should improve long-term tolerability because the immune system is less likely to generate anti-drug antibodies against a fully human (not foreign) molecule.

## The acquisition chain

D2E7 was developed and manufactured at BASF Bioresearch Corporation, then taken through clinical development by BASF's pharmaceutical subsidiary, Knoll Pharmaceuticals. In late 2000, Abbott Laboratories agreed to acquire Knoll Pharmaceuticals from BASF for approximately $6.9B in cash. Abbott, a diversified healthcare company based in Chicago, wasn't primarily a pharmaceutical company at the time. It made

diagnostic equipment, nutritional products (like Ensure and Similac), and medical devices. But D2E7 was promising enough to justify the entire acquisition. This is a pattern that repeats constantly in the industry -- a large company acquires a smaller company primarily for one promising drug or platform.

On December 31, 2002, the FDA approved adalimumab under the brand name Humira (short for Human Monoclonal Antibody In Rheumatoid Arthritis) for the treatment of moderately to severely active rheumatoid arthritis in adults. It was the first fully human monoclonal antibody to receive FDA marketing approval.

In its first full year on the market (2003), Humira generated $280M in revenue and was being used in 38 countries. Impressive by any measure, but nobody at the time would've predicted what came next. For context, when Abbott acquired Knoll for $6.9B, analysts valued D2E7 as a potential $1-2B per year drug at peak, which would've made the acquisition a good deal but not a legendary one (more on this later).

Eventually, in 2013, Abbott spun off its pharmaceutical division as a new company called AbbVie. By then, the drug had moved through five organizations across three countries. That's the system at work.

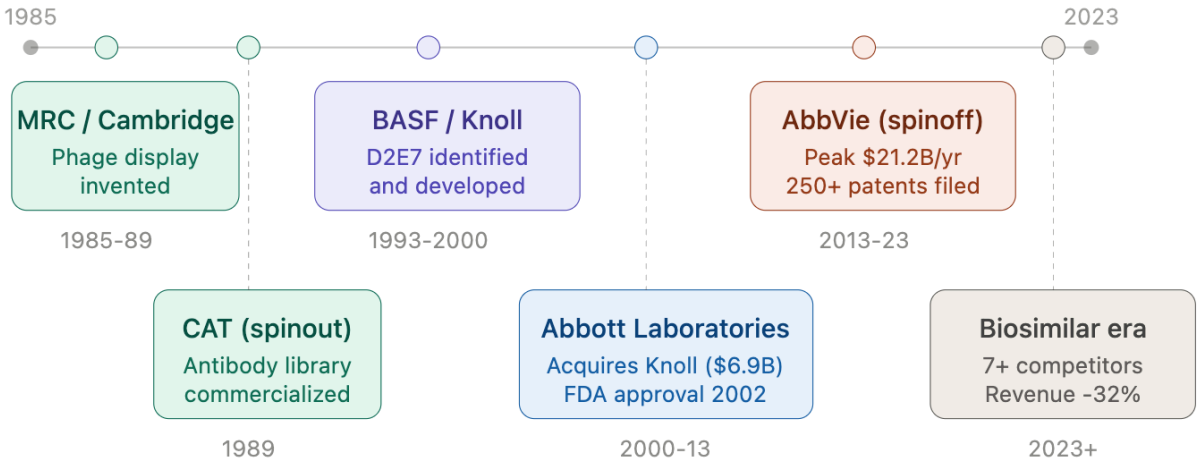

# Part 2 -- getting through the clinic

## The gauntlet

Before any drug can be sold, it must survive the clinical trial process. This is the most expensive, time-consuming, and failure-prone stage in the entire system. Understanding it is essential to understanding why drugs cost what they cost and why the industry is structured the way it is.

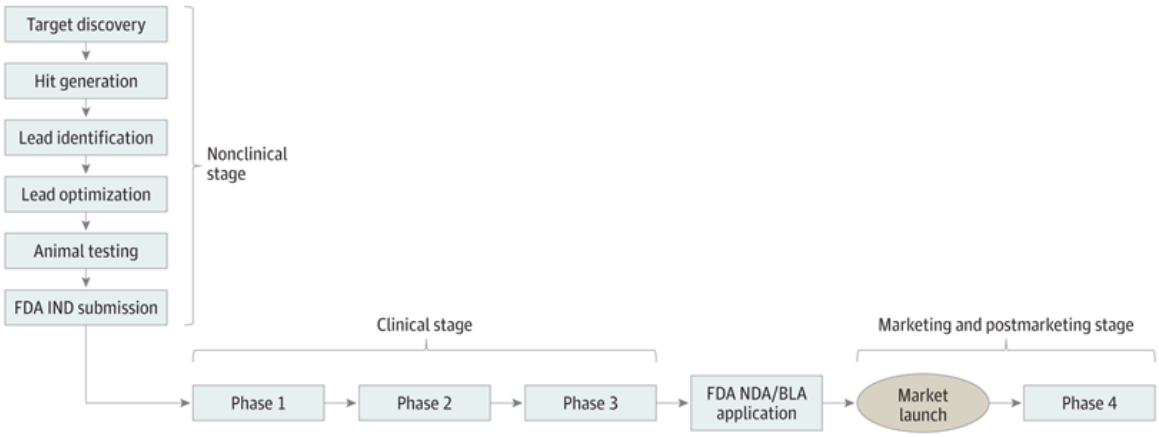


Before human trials even begin, a drug must go through preclinical development: laboratory experiments and animal studies that establish the drug's basic safety profile, mechanism of action, pharmacokinetics (how it's absorbed, distributed, metabolized, and excreted), and optimal dosing range. This stage typically takes 3-5 years and requires the company to file an Investigational New Drug (IND) application with the FDA before human testing. The IND includes preclinical data plus a detailed plan for the proposed clinical trials. The FDA has 30 days to review and either allow the trial to proceed or place a clinical hold. Once the IND is cleared, the drug enters the three phases of clinical testing in humans:

- *Phase 1* tests safety in a small group of healthy volunteers or patients (typically <100) throughout several months to determine safe dosage ranges, characterize side effects, and understand the drug's pharmacokinetics in humans. For oncology drugs, Phase 1 trials are often conducted in patients rather than healthy volunteers (for ethical reasons). Most drugs pass Phase 1 because it's primarily a safety screen, not an efficacy test.
- *Phase 2* tests the drug's efficacy in a larger group of patients with the target disease (typically 100-300) throughout 1-2 years. It's the first test of whether the drug actually works in humans. Phase 2 is divided into Phase 2a (dose-finding) and Phase 2b (efficacy). The failure rate here is roughly 50-60%, which makes it the single biggest bottleneck in drug development. Many drugs that appear promising in preclinical models simply don't translate to human biology.
- *Phase 3* tests efficacy and monitors side effects in an even larger patient population (typically 1,000-3,000+), usually in randomized, double-blind (neither patients nor doctors know who's getting treatment and who isn't), placebo-controlled trials that compare the drug to either a placebo or the current standard of care. These are very expensive (often $100-300M for a single trial), take 2-4 years, and require global recruitment across dozens of clinical sites. Phase 3 trials must demonstrate a statistically significant and clinically meaningful benefit to support regulatory filing. They're also the primary source of safety data for the drug's label (what the drug can be prescribed for).

The overall probability of a drug entering Phase 1 and eventually receiving FDA approval is roughly **10-15%**. A widely cited study in Biostatistics (Oxford) estimated an overall probability of approval of about 13.8%; an MIT analysis arrived at nearly 14%. Another study in Clinical Pharmacology & Therapeutics found that approval rates vary drastically based on factors like drug target, modality, and indication.

What drives these failure rates? In most cases, it's the fundamental challenge that a molecule that binds a target in a test tube or works in a mouse simply doesn't produce the desired effect in humans. The reasons for this are varied -- sometimes the target isn't as important in human disease as animal models suggested; or the

drug hits the target but compensatory pathways in human biology route around the blockade; or the drug works but causes unacceptable side effects that only become apparent in a larger, more diverse patient population; or the drug works beautifully in a subset of patients but not in the broader population, and the trial wasn't designed to detect that subgroup. This is why the clinical trial process is so rigorous and why shortcuts are so dangerous. You genuinely don't know whether a drug works until you test it in enough humans, under controlled conditions, for long enough to see both benefits and harms.

As you can imagine, this gauntlet is expensive. Estimates of the average cost to develop a single approved drug range from $1-3B, depending on methodology. A JAMA study (2000-2018) found a mean development cost for one drug is over $900M. A RAND analysis confirmed that the median is around $708M but the average is much higher ($1.3B) due to outliers. Critically, these figures include the cost of all the drugs that failed along the way. You aren't just paying for the one drug that works; you're also amortizing the cost of the nine or so that didn't.

What's surprising is that it's getting worse, not better. A phenomenon known as Eroom's Law (Moore's Law spelled backwards) describes the observation that the number of new drugs approved per billion dollars of R&D spending has halved roughly every nine years since 1950, falling around 80-fold in inflation-adjusted terms. Despite massive advances in genomics, high-throughput screening, and computational tools, the cost of producing a new drug has relentlessly climbed. The reasons include what the paper's authors call the 'better than the Beatles' problem (every new drug must beat a growing back-catalog of cheap generics), increasingly cautious regulators, and a tendency to overestimate how much brute-force technology can overcome biological complexity.

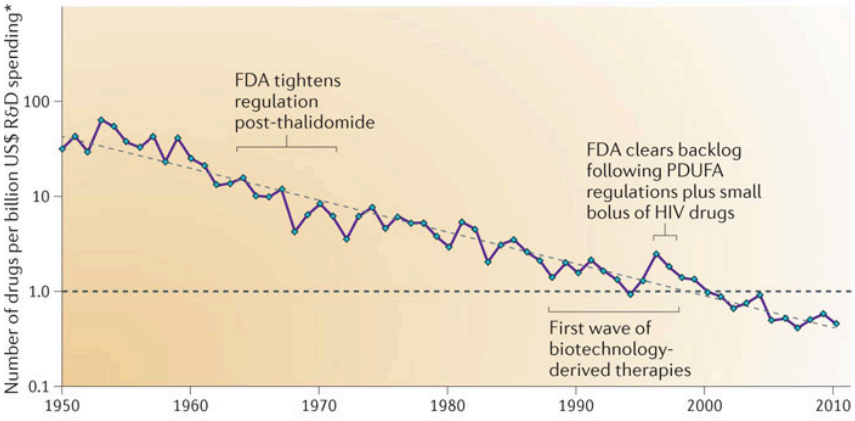


So… can technology reverse this by compressing the timeline or widening the funnel? AI is starting to play a role in the earlier stages: virtual drug screens, clinical trial protocol design, patient selection criteria, safety signal identification, and life-cycle management (covered in the next section). But my honest assessment is that the fundamental bottleneck is biological, not computational. The reason Phase 2 trials fail 50-60% of the time isn't because we designed the trials poorly; it's because the drug didn't work in humans the way it worked in mice. AI can improve the quality of hypotheses entering the clinic (better target selection, better molecule design), and that may modestly improve overall success rates over time. But the funnel is narrow for a reason and, in my opinion, making it wider while maintaining safety and rigor is one of the most impactful problems to work on in this century.

## Humira's label expansion

Humira was initially approved for rheumatoid arthritis, but its mechanism of action (blocking TNF-alpha, a general inflammatory signal) meant it could theoretically work in many inflammatory and autoimmune conditions. TNF-alpha isn't specific to the joints; it drives inflammation throughout the body. After the initial RA approval, Abbott (and later AbbVie) systematically ran new clinical trials to expand Humira's label to additional indications.

Over the following years, Humira received FDA approval for psoriatic arthritis (2005), ankylosing spondylitis (2006), Crohn's disease (2007), plaque psoriasis (2008), juvenile idiopathic arthritis (2008), ulcerative colitis (2012), hidradenitis suppurativa (2015), uveitis (2016), and more. Each new indication required its own clinical trial program and separate regulatory submission, though Phase 1 safety studies could be skipped since the drug's safety profile was already established.

By broadening Humira's addressable patient population from millions of RA patients to tens of millions of people across multiple autoimmune diseases, they built the drug into a commercial juggernaut that would peak at $21.2B in annual revenue in 2022. The fact that Humira eventually peaked at more than 10x the most optimistic early projections tells you something important about drug development: even the experts systematically underestimate the commercial potential of drugs with broad mechanisms of action. When a drug works through a pathway as fundamental as TNF-alpha, the addressable market is much larger than any initial label suggests through successive label expansions, because the same biology drives disease across many organs and conditions.

This strategy (called life-cycle management) is one of the most reliable playbooks in pharma. Find a drug with a broad mechanism, prove it works in one disease, then systematically expand into every adjacent condition where the biology overlaps. A more recent example of this, which we'll return to later, is the GLP-1 receptor agonist class. Drugs like Ozempic were initially approved for type 2 diabetes and later for obesity, with additional indications still being explored. The approach is efficient because the drug's safety profile is already well-known from the initial approval and post-market surveillance, so the incremental risk and cost of each new indication trial is lower than starting from scratch with a new molecule.

## The gatekeeper

The Food and Drug Administration (FDA) is the regulatory body that approves (or rejects) drugs for the US market, and its decisions reverberate globally. The FDA's goal is simple: to ensure that approved drugs are safe and effective. The average FDA review time for a new drug application (NDA, filed after Phase 3) is about 10 months, though several expedited pathways exist for drugs that address particularly urgent medical needs:

- *Breakthrough Therapy Designation* provides intensive FDA guidance and communication for drugs that show substantial improvement over existing treatments.
- *Accelerated Approval* allows drugs to be approved based on a "surrogate endpoint" (a lab measurement or physical sign that's reasonably likely to predict clinical benefit) rather than waiting for definitive clinical outcome data, with confirmatory trials required after approval.
- *Fast Track Designation* facilitates more frequent meetings with the FDA and rolling review of the application.
- *Priority Review*, granted for drugs treating serious conditions with unmet need, shortens NDA review time to about 6 months.

The FDA's role extends beyond simply approving or rejecting drugs. It also sets the evidentiary standards that shape the entire clinical development process. A company's meetings with the FDA before designing a Phase 3 trial (called End-of-Phase-2 meetings) are very consequential. The FDA provides guidance on trial design -- how many patients to enroll, what primary endpoint to use, what statistical threshold to meet, and what safety monitoring is required (all of which influence how expensive and time-consuming the trial will be). A Phase 3 trial that the FDA wants to see run in 3,000 patients across 200 sites for 18 months is a fundamentally different economic proposition than one requiring 500 patients across 20 sites for 6 months. The FDA's guidance isn't arbitrary; it's calibrated to the disease severity, the treatment landscape, and the strength of the Phase 2 data. But it's also the single biggest driver of late-stage development costs, which is why the regulatory interaction is so strategically important.

You may be wondering… how do companies run clinical trials in practice? Most biotechs don't run their own trials operationally. Instead, they work with contract research organizations (CROs) like IQVIA, Covance (now part of LabCorp), or Charles River Laboratories. CROs handle site management, patient recruitment, data collection, regulatory submissions, biostatistics, and pharmacovigilance. Similarly, contract development and manufacturing organizations (CDMOs) like Catalent, Lonza, and WuXi Biologics handle the actual manufacturing of drugs, from small batches for clinical trials to commercial-scale production. Given how expensive manufacturing can be, most companies outsource to CDMOs and only consider building their own facilities if a drug reaches blockbuster status. This outsourced infrastructure ecosystem enables a biotech company with 50 employees to run a global Phase 3 trial.

# Part 3 -- biotech economics

## How venture capital fuels biotech

As we've discussed, drug development is *very* expensive and takes a *very* long time. A typical biotech startup will spend 10-15 years and over a billion dollars before generating any revenue, if it ever does. This creates a financing problem that's fundamentally different from software or consumer tech, where you can ship a product in months and iterate based on user feedback. In biotech, you can't iterate on a molecule mid-trial. You design, test, wait, and learn whether it works only after years and hundreds of millions of dollars.

Venture capital solves this problem, but the structure looks different from tech VC. Biotech VCs index very highly on the team, particularly at the early stages when there's minimal data and the bet is fundamentally on people. A great team with a mediocre asset will often outperform a mediocre team with a great asset because drug development requires navigating so many unpredictable challenges that adaptability and judgment matter enormously.

Each funding round typically corresponds to specific scientific or clinical milestones, and the capital raised is sized to get the company to the next value-inflecting data point(s):

- *Seed / Series A ($5-80M):* The company is founded around a discovery (usually licensed through the university's TTO, as discussed). The money funds initial research, proof-of-concept experiments, and potentially IND-enabling studies (the preclinical work needed to file with the FDA). The team is typically small: a few scientists and a C-suite (CEO, CSO, maybe CMO). In 2025, Series A biotech valuations averaged around $80M pre-money, which is high by historical standards and reflects the capital-intensity of the space.
- *Series B ($50-200M+):* Funds somewhere around Phase 1 and Phase 2 clinical trials. By this point, there's usually some human data (at least early safety data). The risk is still very high (remember, 50-60% of drugs fail in Phase 2), but the potential value has increased. Investors at this stage are typically specialized healthcare VCs who can evaluate clinical data. Median Series B biotech valuations reached $157.5M in Q3 2023, far higher than any other sector.
- *Series C and beyond ($100M-500M+):* Funds Phase 3 clinical trials, which are the most expensive. At this stage, the company may also be preparing manufacturing infrastructure (often through CDMOs) and building a commercial team in anticipation of launch. Some companies raise crossover rounds from a mix of VC and public market investors as a bridge to an IPO.

The scale of biotech VC is significant. In 2024, biotech venture funding reached $26B across 416 rounds. Health and biotech startups dominated US Series A funding in 2024. Contextually, the failure rate in biotech is much higher than in tech, so the winners need to return much more to compensate for all the losses. A single successful drug can return an entire fund. Most portfolio companies return nothing.

Biotech VC, just like other industries, goes through broader and industry-specific economic cycles. The VC model is fundamentally about exits, either selling the company (M&A) or taking it public (IPO). Investors are constantly mapping the milestones between now and the next financing round, and sometimes mapping all the way to the eventual liquidity event, which is when the money returns into the system to be deployed into future companies (this is a healthy cycle).

This creates a cascading dynamic that's hard to appreciate without zooming out and looking at the cyclical nature of this business. When the public markets are open and receptive to biotech IPOs (as they were in 2020-2021), late-stage investors are aggressive, which makes mid-stage investors more comfortable, which makes early-stage investors willing to take bigger swings. The whole system runs hot. But when the IPO window closes (as it did in 2022-2023), that confidence collapses in reverse. Late-stage crossover investors pull back. Series C rounds become harder to close. Series B investors demand more clinical data before committing. This trickles all the way down to seed investors, who now want to see more preclinical evidence before writing a check (often at worse terms). Many companies that raised at inflated valuations during the 2020-2021 boom found themselves raising down rounds in 2023-2024, reflecting the new reality. This isn't unique to biotech (tech experienced the same cycle), but the consequences are more severe in biotech

because you can't simply cut costs and reach profitability. The molecule doesn't care about your runway -- it needs the clinical trial it needs, and that trial costs what it costs.

The structure of biotech VC firms also differs from tech VC. Firms like Flagship Pioneering (which created Moderna), ARCH Venture Partners, and The Column Group are deeply involved in company creation -- identifying scientific opportunities, recruiting founding teams, licensing IP from universities, and serving as acting executives in the early stages. This venture creation model is more common in biotech than in tech and reflects the specialized expertise required to evaluate drug programs.

## The biotech founder's journey

There's a stakeholder in this system who often gets overlooked -- the founder. Building a biotech company is one of the most uniquely demanding forms of entrepreneurship that exists.

The typical founding story begins with a scientist. A professor at a research university makes a discovery with therapeutic potential. Maybe they identified a novel drug target, developed a new screening platform, or engineered a molecule with unusual properties. The university's TTO files a patent. A VC firm approaches the scientist, or the scientist approaches a VC firm, and a company is formed around the discovery. The scientist may become the chief scientific officer or (less commonly and more ambitiously) the CEO. Often, a professional biotech CEO is recruited to run the company while the scientist returns to their lab or takes an advisory role.

What makes biotech founding uniquely difficult is the combination of timeline, cost, uncertainty, and stakes. Your product might take 10-15 years to reach the market and will cost around a billion dollars with roughly 90% chance of failing entirely. If it fails, it's not necessarily because of bad product-market fit or poor user experience; it's because your hypothesis didn't translate to human biology (something nearly impossible to predict in advance). And unlike a software startup where you can pivot, you can't pivot a Phase 2 clinical trial.

You designed the molecule, you enrolled the patients, you ran the study, and now you wait for the data to come back. If it's negative, you may've just spent $100M to learn that your drug doesn't work. There's no iteration loop -- there's only the data readout. That said, there are ways to diversify: building platforms that yield multiple drug candidates, taking multiple shots on goal in parallel. But in practice, most biotech companies rise or fall on their lead asset.

On top of this, biotech founders face relentless dilution. By the time a drug reaches approval (after Seed, Series A, B, C, maybe a crossover round, maybe an IPO), the original founders and early employees may own <5% of the company. This is dramatically lower than in tech startups, where founders routinely retain over 20% at IPO. The dilution reflects the capital intensity of the business -- you need more money, which means more equity rounds, which means more dilution. The founder's financial reward, while potentially significant in absolute terms if the drug succeeds, comes from a much smaller slice of a much larger (and uncertain) pie.

And yet, people do it because drug development is one of the most beautiful orchestrations of effort in modern science. A biotech company at its best is a small team of deeply talented people (biologists, chemists, clinicians, regulatory strategists, manufacturing experts) all pointed at the same impossibly hard problem: can we turn this molecular insight into something that helps a patient? The founder's job is to recruit those people, raise the capital to fund their work, set the scientific strategy, manage the board, navigate regulatory interactions, and keep the whole operation moving forward through years of uncertainty. The best biotech founders aren't just good scientists or good business operators; they're people who can do both at an elite level and switch between them depending on the audience and the moment.

The returns, when they come, aren't just financial. A biotech founder who successfully develops a drug has contributed something that'll outlive them. The founder's work becomes a permanent contribution to human health. There aren't many jobs where you can say that, with confidence, in every case where you succeed.

## Taking the company public

Many biotech companies go public before they have any revenue, or even an approved product. This is unusual compared to most industries and reflects the enormous capital requirements of late-stage clinical trials. An IPO gives a biotech company access to public market capital and also provides liquidity for early investors.

A biotech IPO looks nothing like a tech IPO. When Airbnb or Snowflake go public, investors can look at revenue growth, user metrics, and unit economics. When a preclinical biotech goes public, investors are buying a thesis -- that this molecule, developed by this team, targeting this disease, has some probability of eventually working. The S-1 filing (the document a company files with the SEC before going public) will contain no revenue, customers, or product. What it'll contain is clinical data (or preclinical data, for the really early companies), a description of the market opportunity, a pipeline chart showing each drug candidate and its development stage, the management team's track record, the IP portfolio, and a risk factors section that can run 30+ pages. Biotech investors are essentially VCs operating in the public markets, which is why the space attracts a specialized group of institutional investors (e.g. Baker Brothers, RA Capital, Perceptive Advisors) who have the scientific expertise to evaluate drug programs the way generalist investors can't.

The biotech IPO market is notoriously cyclical. In 2020 and 2021, the market was on fire: 183 biotechs raised nearly $30B combined, with 2021 alone seeing a record 104 IPOs. Investor enthusiasm, driven partly by the success of mRNA vaccines during COVID-19, led to a wave of preclinical and early-stage companies going public. Then the market collapsed. In 2022, only 21 biotechs went public, raising a mere $1.4B. By 2023, the number recovered slightly to about $2.9B.

What fundamentally changes during these cycles? In short -- risk appetite. During 2020-2021, most companies that IPO'd were preclinical or in Phase 1; by 2023, most had drugs in Phase 2 or later. The public markets had essentially recalibrated to demand more clinical de-risking before they'd invest. Many of the companies that went public during the boom at inflated valuations subsequently traded well below their IPO price, destroying billions of dollars in investor value and reminding everyone that pre-revenue biotech is fundamentally high-risk.

The IPO dynamics matter for the entire ecosystem because of the trickle-down effect I described in the VC section. When the IPO window is open, late-stage VCs are more confident in their ability to exit, so they invest more aggressively. That confidence cascades backward through each funding stage, inflating the whole pipeline with capital availability. When the IPO window closes, the reverse happens: each stage becomes more cautious and data-hungry. Companies formed during the boom find themselves stuck with expensive clinical programs and no clear path to the next round. Some raise down rounds. Some merge in rescue deals. Some simply run out of money and shut down, with their drug candidates returning to the university or being acquired at fire-sale prices. This cyclicality is one of the most important features of the biotech ecosystem.

## The big handoff

The most common exit for a successful biotech company is acquisition by a larger pharmaceutical company. This is the primary mechanism by which big pharma replenishes its pipeline, and the numbers are large. In 2023, biopharma M&A represented approximately $152B in aggregate deal value, the highest since 2019. Notable deals included Pfizer's $43B acquisition of Seagen (an antibody-drug conjugate company), Bristol Myers Squibb's $14B purchase of Karuna Therapeutics (at a 53% premium to the share price, for a schizophrenia drug), Merck's $10.8B acquisition of Prometheus Biosciences, AbbVie's $10.1B purchase of ImmunoGen, and AbbVie's $8.7B acquisition of Cerevel Therapeutics.

2024 was quieter, with aggregate M&A deal value of approximately $48B (down 68% from 2023) and an average deal size of just $2.1B. But the underlying structural trend is clear -- pharma needs biotech more than biotech needs pharma. Large pharmaceutical companies face constant patent cliffs as their blockbuster drugs lose exclusivity, and internal R&D hasn't been productive enough to replace them. Acquiring biotech companies is how pharma fills that gap. This process, too, follows industry cycles that determine the deal flow and burden of proof required for a deal.

The M&A process illustrates the system's incentive alignment. The price is typically set as a premium to the biotech's current stock price (for public companies) or as a multiple of the drug's risk-adjusted net present

value (rNPV, for private companies). The premium reflects the acquirer's confidence that it can execute development more efficiently and commercialize more effectively. The biotech's founders and early investors get liquidity. The pharma company gets a pipeline asset that has been partially de-risked. The entire transaction is driven by complementary capabilities: the biotech can't scale globally, while the pharma company can't discover efficiently. The acquisition is how the system converts one form of comparative advantage into another.

Humira's story illustrates this perfectly. Abbott didn't discover adalimumab. It acquired it by buying Knoll Pharmaceuticals. Then, in 2013, Abbott made a bold structural move -- it spun off its entire pharmaceutical division as a new, standalone company called AbbVie. The logic was focus. AbbVie would be a pure-play pharmaceutical company whose focus was Humira, free to make acquisitions and pipeline investments without being constrained by a parent company's diversified priorities. Abbott retained its diagnostics and medical device businesses. The spinout unlocked enormous shareholder value: AbbVie's market cap eventually exceeded $250B.

## How biotech companies are valued

Valuing a biotech company is fundamentally unique because most have no revenue early on, so their primary asset is probability-weighted future cash flows from drugs in development. The standard method is rNPV. The idea is straightforward. You estimate the potential revenue of a drug if it's successfully approved and commercialized -- based on the size of the patient population, expected market share, pricing, and duration of market exclusivity. Then you discount those future revenues back to present value using a discount rate that reflects the time value of money. Then, critically, you multiply each future cash flow by the *probability* that the drug actually reaches that stage. A preclinical drug might have a 5-10% probability of reaching the market; a Phase 3 drug might have a 50-60% probability. As the drug advances through clinical trials, the probability adjusts upward, and the rNPV increases accordingly. This is why successful clinical data readouts cause such dramatic stock price movements: a positive Phase 2 result can increase a drug's probability-weighted value by 2-5x overnight.

In practice, rNPV models are built in elaborate spreadsheets with dozens of assumptions: patient prevalence and incidence data, diagnosis rates, drug penetration curves, pricing by geography, patent expiration dates, probability of success (PoS) at each stage, cost of goods sold, commercial launch costs, and discount rates that vary by risk tier. The output is a single number that supposedly represents the value of the drug program. The problem, of course, is that every assumption is uncertain. Small changes in key inputs (peak market share going from 20% to 25%, or the probability of Phase 2 success going from 30% to 40%) can change the rNPV by billions. This is why biotech valuation is as much art as science, and why experienced biotech investors spend enormous time debating the biological plausibility of the therapeutic hypothesis, the quality of the clinical data, and the competitive dynamics of the market. The spreadsheet is just the quantification of a judgment call.

Discount rates for early-stage biotech assets can be as high as 50%, reflecting the extreme uncertainty. For late-stage assets approaching launch, discount rates drop to 15-20%. This is why, on paper, the same drug can be worth $100M to a startup and $10B to a pharma company -- the pharma company is assigning a higher PoS (because it has the resources to execute) and a lower discount rate (because it has a diversified pipeline).

# Part 4 -- intellectual property

## How drug patents work

Patents are the economic engine of the industry. Without patent protection, no company would invest a billion dollars developing a drug that a competitor could copy the day it was approved. The basic deal is straightforward -- in exchange for publicly disclosing your invention, the government grants you a 20-year monopoly from the date you file.

In practice, the math isn't that clean. Because patents are often filed during preclinical research, and it takes 10-15 years to bring a drug to market, the actual period of commercial exclusivity is usually around 10 years. The Hatch-Waxman Act of 1984 addresses this by allowing companies to extend their patent term by up to 5 years to recover time lost during FDA review. However, the total patent life after approval can't exceed 14 years. For Humira, the original composition-of-matter patent was filed in the mid-1990s during preclinical development. The drug wasn't approved until late 2002, meaning roughly 7-8 years of the 20-year patent term had passed by the time Humira generated its first dollar. This is the fundamental tension of drug patents -- the clock starts ticking long before the drug makes money.

For biologic drugs like Humira (large, complex molecules produced in living cells, as opposed to small-molecule drugs synthesized through chemistry), there's an additional layer of protection: 12 years of data exclusivity. This means the FDA can't approve biosimilars (the biologic equivalent of a generic) for 12 years after the drug's approval, regardless of patent status. For small-molecule drugs, this data exclusivity period is only 5 years. The longer exclusivity for biologics reflects their complexity. Developing a biosimilar is much more difficult and expensive than developing a generic small-molecule drug, because biologics are made by living cells and even small manufacturing differences can affect the product.

There are several types of drug patents (a single drug can be covered by multiple patents across categories). Understanding them helps explain why patent strategy has become such a sophisticated (and controversial) part of the industry:

- *Composition-of-matter* patents cover the drug molecule itself; these are the strongest and most valuable.
- *Method-of-use* patents cover specific therapeutic applications (treating a particular disease with the drug).
- *Formulation* patents cover specific ways of preparing or delivering the drug (a particular concentration, buffer, or delivery device).
- *Process* patents cover specific manufacturing methods.

## Humira's patent strategy

Humira's original composition-of-matter patent was set to expire in 2016. In a straightforward world, biosimilar competitors would've entered the market around that time, prices would've declined, and the story would've ended there. Instead, AbbVie deployed one of the most aggressive and sophisticated patent strategies in pharmaceutical history.

AbbVie filed over 250 patent applications related to Humira in the US (the exact count varies by source and methodology), approximately 90% of which were filed after the drug's original 2002 approval. Of these, around 150 were granted. These patents covered everything from specific formulations and manufacturing processes to treatment methods for each indication, the citrate-free version of the drug, the pre-filled syringe device, and even specific dosing regimens.

This practice, known as a "patent thicket," is controversial but legal. The basic strategy is to surround the core drug with so many overlapping patents that any biosimilar manufacturer faces a complex legal challenge. Even if a biosimilar company could design around any individual patent, the sheer number of patents means years of litigation. AbbVie used this leverage to negotiate settlements with biosimilar manufacturers, granting them licenses to launch but not until 2023, effectively extending its US monopoly by about six years beyond the

original patent expiration. During those additional years (2016-2022), Humira continued generating $15B+ annually. The patent thicket strategy generated roughly an additional $80-100B that AbbVie wouldn't have received had biosimilar competition begun in 2016.

Whether you view this as brilliant strategy or rent-seeking depends on your perspective. AbbVie would argue that the additional patents reflected genuine innovation: a citrate-free formulation that reduced injection pain, improved manufacturing processes, new treatment regimens optimized for different diseases. Critics argue that many patents were incremental and designed primarily to prevent competition and maintain pricing power. In reality, it's probably both. But the broader point about the ecosystem is that patents aren't just legal protection -- they're the economic asset that makes the entire risk-reward structure of drug development work. Without the promise of market exclusivity, no investor would fund a billion-dollar drug development program. The debate isn't about whether patents should exist (they should), but about how long effective exclusivity should last and what kinds of innovations should be patentable.

And here's the part of the patent story that I think doesn't get enough attention, because it's actually beautiful. When a patent ends, the molecule becomes public. After all the investment, risk-taking, clinical trials, regulatory approvals, and years of commercial exclusivity, the drug enters the public domain. Generic manufacturers (for small molecules) and biosimilar manufacturers (for biologics) can now produce it, leading to price drops and expanded access. A molecule that was conceptualized from a screening library, optimized through rounds of medicinal chemistry, tested in thousands of patients across trials, reviewed by regulatory bodies, manufactured at industrial scale, and prescribed to millions of people… now belongs to everyone. Patients who live a century from now will still benefit from Humira. The system is designed so that innovation gets rewarded temporarily and then shared permanently. That's an elegant deal.

# Part 5 -- the role of big pharma

## What pharma actually does

If biotech startups are the scouts and innovators, big pharmaceutical companies are the ones who scale, manufacture, distribute, and commercialize drugs globally. These capabilities are difficult to replicate, and they justify pharma's role in the ecosystem.

*Manufacturing at scale.* Biologic drugs like Humira are produced in massive bioreactors using living cells. Building a biologics manufacturing facility costs $200M-2B and takes several years. The manufacturing process itself is sensitive -- small changes in temperature, pH, growth media composition, or cell culture conditions can alter the product's structure and potency. This is fundamentally different from manufacturing a small-molecule drug like aspirin, where the chemistry is reproducible and well-understood. It's also one reason biologics are much harder to copy than traditional drugs, and why biosimilars are more expensive and complex to develop than generic pills. The GLP-1 drugs illustrated this challenge vividly -- Novo Nordisk and Eli Lilly both experienced prolonged supply shortages (which persisted for years) because demand for semaglutide and tirzepatide far outstripped their manufacturing capacity. Manufacturing is genuinely hard, capital-intensive, and a competitive advantage for companies with the infrastructure.

*Global regulatory navigation.* Getting a drug approved in the US is just the start. There are separate regulatory agencies in Europe (EMA), Japan (PMDA), China (NMPA), the UK (MHRA), and dozens of other countries, each with their own data requirements, filing procedures, and timelines. Big pharma has teams of regulatory affairs professionals who manage this process across 50+ markets simultaneously. A global launch can take years after the initial US approval.

*Commercial infrastructure.* Selling drugs isn't like selling software. In the US, pharma companies employ thousands of sales representatives who visit physicians and hospitals to educate them about new drugs. They run medical education programs, sponsor continuing medical education for physicians, build relationships with key opinion leaders in each therapeutic area, manage patient support programs (including copay assistance and nurse educator services), and execute direct-to-consumer advertising campaigns (the US and New Zealand are the only countries that allow DTC pharmaceutical advertising). They also negotiate with payers, which is its own complex art.

*Post-market surveillance (Phase 4).* After FDA approval, companies must monitor their drugs for safety signals in real-world use. Clinical trials, even large ones, may not catch rare side effects that only appear in very large patient populations (e.g. 1 in 50,000 patients) or over longer time periods (e.g. after 5 years of use). Pharmacovigilance teams collect and analyze adverse event reports from doctors and patients, and the FDA can require post-market studies or label changes if new safety concerns emerge.

The relationship between big pharma and biotech isn't purely transactional. Many pharma companies maintain extensive partnership networks through licensing deals, research collaborations, option agreements, and co-development arrangements. Roche and Genentech (its biotech subsidiary) are deeply intertwined. Bristol Myers Squibb built its immuno-oncology franchise largely through partnerships and acquisitions (buying Celgene for $74B in 2019). Johnson & Johnson's Janssen unit operates a venture arm and incubator network that collaborates with hundreds of external partners. AstraZeneca rebuilt its oncology pipeline almost entirely through external deals. The point is that big pharma is less a self-contained R&D machine and more a platform for scaling, manufacturing, and commercializing innovations that largely originate elsewhere. Understanding this clarifies why pharma companies are valued the way they are -- their competitive advantage isn't in discovery but in everything that happens after.

## Patent cliffs

The existential challenge facing every drug program is the *patent cliff* -- the moment when a blockbuster drug loses patent exclusivity and generic (or biosimilar) competitors enter the market, causing a dramatic revenue drop. This isn't a theoretical concern. It's the defining strategic pressure of the industry, and every major

pharma company deals with it. Being capital- and time-efficient during trials is critical as it directly influences how much time of exclusivity you have after approval.

The scale of the looming patent cliffs across the industry is staggering. Between 2025 and 2030, drugs generating over $200B in combined annual revenue will face patent expirations or biosimilar competition. This includes Keytruda (~$30B), Opdivo (~$9B), Eliquis (~$13B from BMS's share alone), and others. Each patent cliff forces the originator company to find replacement revenue through a combination of internal R&D, M&A, and life-cycle management of existing assets. Companies that manage this transition successfully survive; if not, they'll see sustained revenue declines, layoffs, and potential acquisition by larger competitors.

U.S. drug revenue by number of years from date of main patent expiration, which is set to Year 0.

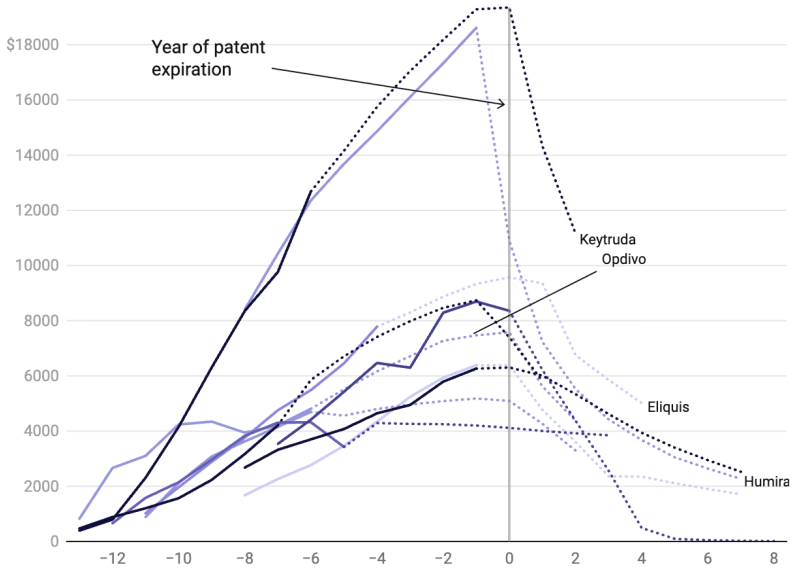


*Solid lines represent actual revenue, through 2021. Dotted lines represent projected revenue. Numbers in millions USD.*

Chart: Julia Himmel / BioPharma Dive • Source: Companies, Evaluate Pharma • Created with Datawrapper

When Humira began losing US exclusivity in 2023, AbbVie faced this challenge. Amgen launched Amjevita, the first Humira biosimilar, at two price points: one priced 55% below Humira's list price, and another priced just 5% below. At least seven more biosimilars entered the market through 2023. As a result, Humira's annual sales fell from a peak of $21.2B in 2022 to $14.4B in 2023, a 32% decline in a single year.

AbbVie had been preparing for this cliff for years by acquiring new drugs through M&A and developing internal candidates. Its newer immunology drugs, Skyrizi (risankizumab, an IL-23 inhibitor) and Rinvoq (upadacitinib, a JAK inhibitor), are designed to eventually replace Humira's revenue. But the transition is painful. It's an endless treadmill for every pharma company, and those that run it best will win.

# Part 6 -- drug pricing

## Who pays for drugs?

In the US, the question of who pays for drugs is more complicated than in almost any other country. Most developed nations have single-payer or regulated healthcare systems where the government negotiates drug prices centrally. The US doesn't (with limited recent exceptions). Instead, the system involves multiple layers of intermediaries between the drug manufacturer and the patient.

The major payer categories are commercial insurance (provided by employers through companies like UnitedHealthcare, Aetna, and Cigna, covering ~150M Americans), Medicare (federal program for people 65+ and certain disabled individuals, covering ~65M Americans), Medicaid (the joint federal-state program for low-income individuals, covering ~90M Americans), and uninsured or self-pay patients (~25M Americans).

## The incentive architects

To understand why pharmacy benefit managers (PBMs) matter, one must understand them through the lens of incentives, because that's what this entire system runs on. PBMs sit between the manufacturer (pharma company) and the insurance company and are one of the most powerful and least understood players in the system. They emerged in the 1960s as claims processors but have evolved into intermediaries that negotiate drug prices with manufacturers, design formularies (list of drugs your insurance covers), manage pharmacy networks, and determine patient cost-sharing.

The market is extremely concentrated: the three largest PBMs (OptumRx, owned by UnitedHealth Group; Express Scripts, owned by Cigna; and CVS Caremark, owned by CVS Health/Aetna) manage roughly 79% of all prescription drug claims in the US. Each of these PBMs is vertically integrated with an insurance company, which creates both efficiency and conflicts of interest.

PBMs are important because they sit where clinical decisions (which drug should a patient take?) intersect with economic decisions (what should a drug cost?), which in turn shapes what doctors prescribe, which in turn shapes what drug companies invest in developing. It's safe to say they're the incentive architects of the prescription drug market. A PBM's decision to place Drug A on Tier 2 and Drug B on Tier 3 propagates through the entire system: manufacturers adjust their pricing and rebate strategies, doctors adjust their prescribing patterns, and patients adjust their adherence. The PBM's incentive is to reduce total drug spending for its client (insurer or employer), and it does this by leveraging its market position to extract rebates from manufacturers. In theory, this creates a powerful cost-containment mechanism. In practice, the rebate system has created complications that are still being sorted out.

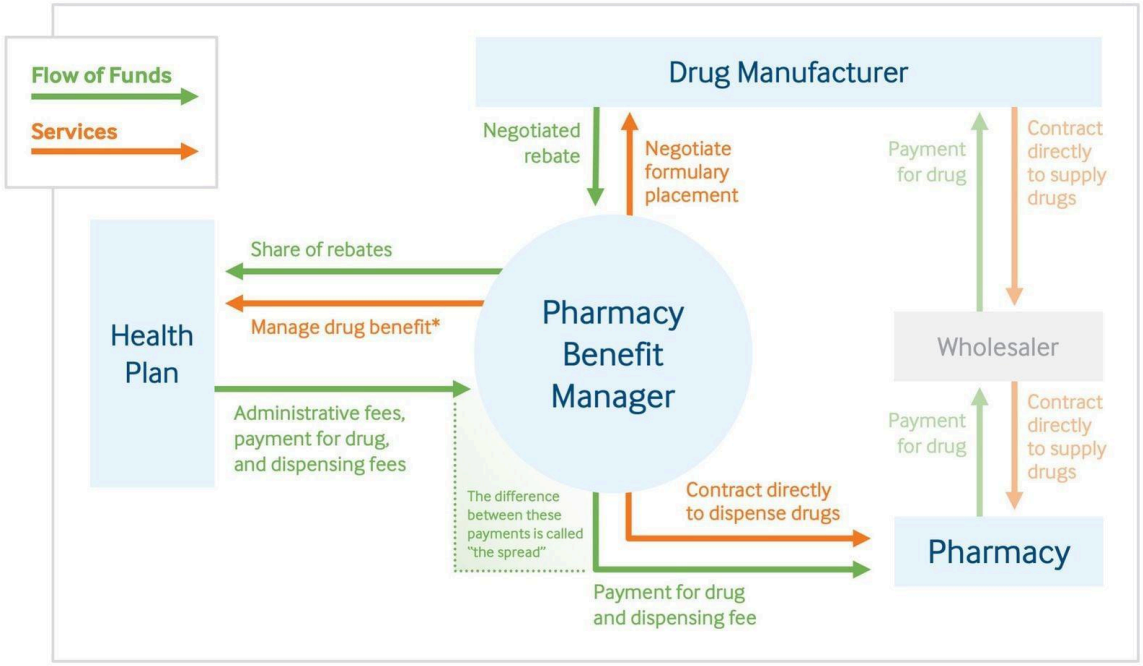


Here's how the money flows… A drug manufacturer sets a "list price" (also called WAC, wholesale acquisition cost). The PBM negotiates a confidential rebate from the manufacturer in exchange for placing the drug on a favorable formulary tier (lower tiers have lower copays, which drives more volume). The insurance company pays something closer to the net price (list price minus rebates). The patient pays a copay or coinsurance at the pharmacy. The PBM retains a portion of the rebates and profits from the spread between what it charges the insurer and what it pays the pharmacy.

Humira nicely illustrates this complexity. Between 2013 and 2020, Humira's list price per dose increased from $1,153 to $2,784 (a 141% increase). But PBM rebates also increased over the same period, from about 2.4% of list price to nearly 35%. So the net price that insurers actually paid increased less dramatically. This dynamic, where list prices rise while rebates also rise, is often called the "gross-to-net bubble." It means the sticker price you see in news headlines isn't the price anyone actually pays, but it creates problems for patients with high-deductible plans or coinsurance percentages, who may pay based on list price rather than net price.

The opacity of the rebate system is one of the most criticized features of US drug pricing. Manufacturers are incentivized to set high list prices because PBMs favor drugs with large rebates (a drug with a $3,000 list price and a 40% rebate generates a $1,200 rebate for the PBM, while a drug with a $1,500 list price and a 10% rebate generates only $150). PBMs prefer drugs with large rebates because their compensation is often tied to the absolute dollar value negotiated. Insurers benefit from lower net prices but may not pass the full savings to patients. And patients with percentage-based coinsurance pay based on the list price, not the net price. This is why a drug can have a list price of $80,000 per year, a net price of $50,000, and still leave a patient owing $20,000 out of pocket. PBMs serve a real function (aggregating bargaining power, managing formularies, processing claims at massive scale), but the incentive misalignments have real consequences for patients.

The entry of biosimilars is supposed to bring prices down significantly. Blue Shield of California announced it would purchase a Humira biosimilar at $525 per monthly dose versus the roughly $2,100 net price of branded Humira. Kaiser Permanente projected $300M in annual savings after switching most patients to biosimilars. Express Scripts announced it would remove branded Humira from its largest commercial formulary entirely, in favor of biosimilars. But the transition has been slower than expected, partly because of complex PBM rebate structures that can make a discounted branded product more profitable for the PBM than a cheaper biosimilar.

## The pricing debate

Drug pricing in the US is a topic that generates lots of political and public attention, and for good reason. Americans pay significantly more for prescription drugs than people in any other developed country. The

reasons are structural: the US didn't historically have centralized price negotiation for most drugs; the PBM system creates incentive misalignments; the 12-year exclusivity period for biologics, combined with patent thickets, means branded drugs face no competition for extended periods; and the US market effectively subsidizes drug prices in other countries, where governments negotiate much lower prices.

The Inflation Reduction Act (IRA) of 2022 represented a significant, if limited, shift. For the first time, it gave Medicare the authority to negotiate prices directly with drug manufacturers for a small number of high-expenditure drugs. The first round covers 10 drugs with negotiated prices taking effect in 2026, with 15 additional drugs each for 2027 and 2028, and 20 additional drugs per year starting in 2029. The negotiated prices represent minimum discounts of 38% off 2023 list prices and are estimated to save Medicare $6B annually and reduce out-of-pocket costs for beneficiaries by $1.5B per year.

The IRA also introduced other important provisions. Medicare beneficiaries now have an annual out-of-pocket spending cap of $2,000 (starting in 2025), which protects the patients who need the most expensive drugs from catastrophic costs. Insulin copays under Medicare are capped at $35 per month. And pharmaceutical companies must pay rebates to Medicare if they raise drug prices faster than inflation. Together, these provisions represent the most significant change to US drug pricing in decades. They don't fix the system, but they begin to address some of the most egregious misalignments between drug costs and patient burden.

That being said, the high prices also fund the R&D that produces new drugs. The US pharmaceutical industry spent roughly $104B on R&D in 2024. This investment enables the 10-15 year, billion-dollar journey from lab bench to approved drug. The tension between access (lower prices, more patients treated today) and innovation (higher returns funding R&D for future drugs) is a central policy debate in the industry, and there's no easy answer. People of good faith disagree about where the line should be drawn.

# Part 7 -- doctors and patients

## How doctors choose treatments

In the middle of this entire system sits the physician. Doctors are the decision-makers who translate clinical evidence into actual prescriptions. A rheumatologist deciding whether to prescribe Humira for a patient with rheumatoid arthritis is operating at the intersection of clinical evidence, professional guidelines, insurance coverage, patient preference, and personal experience.

The process is shaped by several factors: clinical trial data and peer-reviewed literature form the evidence base; professional societies (like the American College of Rheumatology) publish treatment guidelines that rank therapies by recommendation strength -- for many conditions, biologics like Humira aren't first-line therapy; patients typically try less expensive drugs first (methotrexate for RA, for example) and only escalate to biologics if those fail -- this is called "step therapy" and is often required by insurance companies as a cost-containment measure. The physician navigates what the evidence says, what insurance covers, and what the patient can afford and tolerate. Real-world prescribing often diverges from clinical trial evidence for practical reasons: a drug that showed superior efficacy in a randomized trial might be impractical for a patient who can't travel to an infusion center every two weeks, or unaffordable for a patient whose insurance requires a $500 monthly copay. The physician becomes an optimizer across multiple dimensions (efficacy, safety, convenience, cost, patient preference), and different physicians weigh these dimensions differently based on their training, practice setting, and patient population. This is an optimization exercise that clinical decision support tools, increasingly powered by AI, are beginning to help with.

Pharma invests heavily in medical affairs teams and medical science liaisons who engage with physicians scientifically, sharing clinical data, answering questions, and discussing treatment strategies. This is distinct from traditional sales (though the boundary can blur). Companies also sponsor continuing medical education programs, fund clinical research at academic medical centers, and work with key opinion leaders who influence prescribing behavior through their publications, conference presentations, and participation on guideline committees. These relationships are an important (and scrutinized) dynamic in pharma commercialization.

## Patients at the center of it all

For the patient, the system is often bewildering. Consider someone diagnosed with moderate-to-severe plaque psoriasis who's prescribed Humira. Their experience might look like this… The doctor writes a prescription. But the insurance company requires prior authorization (proof that the patient has tried and failed cheaper treatments). The doctor's office submits paperwork documenting previous treatments. The insurer reviews and approves (or denies, requiring an appeal, which can take weeks). If approved, the patient may receive the drug through a specialty pharmacy (biologics aren't typically dispensed at your local CVS). AbbVie's copay assistance program may reduce the patient's out-of-pocket cost to $5 per month. The patient self-injects every two weeks using a pre-filled pen and follows up with their doctor regularly to monitor efficacy and side effects.

That patient may never think about the Cambridge lab where Gregory Winter pioneered antibody phage display, the German chemical company that funded the early research, the 250+ patents that kept the price high, or the PBM that negotiated the rebate that determined their copay. But all of those actors shaped the product that's now helping their skin clear. This is why understanding the system matters -- every decision made upstream eventually impacts patients.

# Part 8 -- inflection points and value creation

One framework I find particularly useful for understanding biotech is thinking about inflection points -- the specific moments where a drug program's PoS (and therefore its value) changes dramatically. These are the moments where the most value is created or destroyed, and understanding them explains much of the behavior in the industry. The key ones, roughly in chronological order, through a drug's life:

- *Target validation.* Preclinical data convincingly shows that modulating a specific biological target has the desired therapeutic effect in animal models. This is where most academic discoveries stall (animal models often don't predict human outcomes), but the ones that clear this bar become fundable biotech companies. The company might be valued at $10-50M.
- *IND filing and Phase 1 initiation.* The FDA allows human testing to begin. This is a regulatory milestone that significantly de-risks the program (it means the FDA has reviewed the preclinical data and found it adequate to support human testing). The company might raise a $50-100M Series A or B.
- *Phase 2 proof of concept.* This is the single most important data readout in a drug's development. Phase 2 is where you learn whether the drug actually works in humans with the target disease. If the data is positive, PoS increases from around 20% to around 60%. This is often the trigger for M&A interest from pharma companies because the risk-reward profile has shifted dramatically. It's also the point of maximum information asymmetry -- insiders know the data before the market does, which is why Phase 2 data readouts can cause 50%+ stock price swings in either direction.
- *Phase 3 success.* Large-scale trials confirm efficacy and safety. The drug is now highly likely to be approved and the remaining risk is mostly commercial. The company might be valued at a few billion at this stage, depending on the market opportunity.
- *FDA approval.* The drug can now be sold. But approval alone doesn't guarantee commercial success; the drug still needs to be adopted by doctors, covered by insurance, priced correctly, and manufactured reliably at scale.
- *Label expansion.* Approval in additional indications can multiply the drug's commercial potential many times over. Humira's expansion from one indication to 10+ is an extreme example. Each new indication adds a new revenue stream to the same underlying asset.
- *Patent cliff.* When exclusivity is lost, revenue typically drops 50%+ within a few years as generics or biosimilars enter the market. For pharma companies, this is an existential event that must be managed through pipeline replenishment, M&A, or life-cycle management.

Understanding these inflection points explains a lot of behavior. Pharma companies pay huge premiums to acquire biotechs right after positive Phase 2 data because the risk-reward profile just shifted dramatically. Biotech stocks swing wildly on data readouts because a Phase 2 miss can reduce a company's value (rNPV) by 80% overnight. Companies file hundreds of patents because every year of extended exclusivity before the patent cliff can be worth billions. The list goes on.

Let's look at a hypothetical example… Imagine a biotech company developing a new biologic for moderate-to-severe Crohn's disease, which affects roughly 500,000 patients in the US. If the drug is priced at $50,000 per patient per year and captures 20% market share at peak, that is $5B in annual peak revenue. Apply a PoS -- in Phase 1, the rNPV might be $5B times 14% (probability of eventual approval) times a discount factor (say, 0.3 for early stage), which gives you a risk-adjusted value of roughly $210M. After a positive Phase 2 result, PoS jumps to 55%, the discount factor improves to 0.5, and the rNPV is now $1.4B. That's nearly a 7x increase in value from a single data readout. This is why Phase 2 data events are the most closely watched moments in biotech. It's also why VC investors time their entry before these inflection points and why acquirers time their bids to be just after them. The entire capital structure of the industry is organized around these probability-adjusted value jumps.

## Part 9 -- mapping the incentives machine

Now that we've followed Humira's story through every stage, let me step back and describe the ecosystem as a whole. Think of it as a machine with eight primary parts (stakeholders), each with distinct incentives, all connected by the push and pull of science, capital, and risk.

- *Academic scientists and research institutions* are incentivized by publications, grants, tenure, and occasionally equity in spinout companies. Their interest is discovery. The system converts their discoveries into drugs through licensing and spinouts. Government funding (primarily NIH in the US) provides investment that makes early-stage discovery possible.
- *Biotech founders, entrepreneurs, and startups* are incentivized by a combination of equity value creation, scientific impact, and organizational building. Their interest is translation -- taking a scientific discovery, orchestrating the extraordinary coordination of talent, capital, and regulatory strategy needed to turn it into a medicine, and either selling the company (to pharma via M&A) or taking it public (via IPO). Founders' equity gets diluted (on top of high failure rates), but a successful outcome can still be life-changing financially while also being deeply meaningful.
- *Venture capitalists* are incentivized by fund returns. They absorb risk but expect outsized returns from the winners (power law). A single successful drug can return an entire fund. They provide capital that bridges academic discovery to clinical proof and bring operational expertise to help companies navigate the development process.
- *Big pharma* is incentivized by revenue and shareholder returns. Their core competency is clinical development at scale, manufacturing, regulatory navigation, global commercialization, and life-cycle management. They need a constant stream of drugs to replace those losing patent protection, and they acquire that stream through a combination of internal R&D, licensing deals, partnerships, and M&A.
- *Regulators (FDA, EMA, etc.)* are incentivized by public safety. They must balance two competing risks: approving an unsafe drug (Type 1 error) and delaying or rejecting an effective one (Type 2 error). Their role is to ensure that approved drugs are safe and effective, and they set the evidentiary standards that the entire clinical development process is designed to meet.
- *Payers (insurance companies, PBMs, Medicare, Medicaid)* are incentivized to minimize costs while maintaining adequate coverage for their members. They act as the economic check on the system, negotiating prices, managing formularies, and determining which drugs patients can access and at what cost.
- *Doctors* are incentivized by patient outcomes, professional reputation, evidence-based practice, and income. They translate clinical evidence and insurance coverage into prescriptions.
- *Patients* want effective treatments that they can access and afford. They're the ultimate beneficiaries (and sometimes victims) of every decision made upstream.

The system works because each group is necessary and interdependent. Pharma can't launch a drug without FDA approval. Doctors can't prescribe a drug without clinical evidence. Biotech can't develop a drug without VC funding. VC can't get returns without pharma acquisitions or public markets. Academia can't translate discoveries without biotech or pharma partners. Payers can't manage costs without biosimilar competition and formulary leverage. And none of it matters if patients can't access treatments. This interdependence is what makes it a *machine*.

The beauty of this arrangement is in the alignment. Each player pursuing their own interest contributes to the shared goal of producing effective medicines. This isn't altruism -- it's brilliant incentive design. The scientist who publishes a breakthrough paper and the VC who funds the company that licenses it aren't collaborating out of goodwill. They are each responding to their own incentive structure. But those individual incentives converge on the same outcome. The scientist gets cited. The VC gets returns. The founder builds an organization. The pharma company fills a pipeline gap. The doctor gets a new treatment. The patient gets relief. This is what I mean when I say the system is *beautiful*: not that every outcome is just or every price is fair, but that the incentive architecture channels an extraordinary amount of human ambition and capital toward a genuinely valuable end.

When it works well, the result is a drug like Humira that transforms the lives of millions of patients with debilitating autoimmune diseases. When it doesn't, drugs get stuck between academic discovery and

commercial development, or reach patients at prices they can’t afford, or get shelved inside a large corporation after an acquisition. The system isn’t perfect -- but understanding where it succeeds and fails requires understanding the incentive structure, which is the purpose of this writeup.

# Part 10 -- two stories worth knowing

## The weight loss revolution

If Humira is the story of one drug's dominance over two decades, the GLP-1 drugs are the story of what happens when an entire market explodes at once. GLP-1 receptor agonists (the class that includes Ozempic, Wegovy, Mounjaro, and Zepbound) were originally developed for type 2 diabetes, based on the gut hormone GLP-1 first described in the early 1980s.

The key scientific breakthrough came from Novo Nordisk, who figured out how to make the GLP-1 peptide last long enough in the body to be a practical medication. The natural hormone degrades within minutes; their engineering extended its half-life to a week. This led to semaglutide, approved for diabetes (as Ozempic) in 2017 and for obesity (as Wegovy) in June 2021. In the landmark STEP 1 trial, semaglutide produced roughly 15% body weight loss over 68 weeks -- unprecedented numbers for a weight-loss drug.

Then Eli Lilly entered with tirzepatide, a dual GLP-1/GIP receptor agonist marketed as Mounjaro (diabetes, 2022) and Zepbound (obesity, 2023). Tirzepatide showed even greater weight loss of 20-22%, and in a head-to-head trial demonstrated clear superiority over semaglutide. In the first nine months of 2025, Lilly's Mounjaro and Zepbound generated $39.5B in revenue, surpassing Keytruda (which we'll discuss next) as the world's best-selling medicine and helping Lilly become the first pharmaceutical company to exceed a $1T market capitalization.

In December 2025, the FDA approved the Wegovy pill, the first oral GLP-1 for chronic weight management, showing about 16.6% weight loss at 64 weeks -- comparable to the injection. This was a major milestone: oral GLP-1s eliminate the need for weekly injections and could dramatically broaden the addressable patient population.

What makes the GLP-1 story fascinating from an ecosystem perspective is the swarming effect. Once the clinical and commercial viability of GLP-1s for obesity was proven, virtually every major pharma company piled in. By 2025, there were over 160 obesity drugs in development covering 68 mechanisms of action. The competition is fierce: Eli Lilly's retatrutide (a triple-agonist targeting GLP-1, GIP, and glucagon receptors) showed 28.7% weight loss at 68 weeks in a Phase 3 trial (TRIUMPH-4, 2025), the highest ever reported for a pharmaceutical agent. Novo Nordisk's CagriSema, combining semaglutide with a novel amylin analog, lost to tirzepatide in a head-to-head Phase 3 trial in early 2026, a significant setback that sent Novo's stock down over 20%. This swarming is the system working as designed. One company proves a concept, competitors pile in with improved versions, prices eventually drop, and patients get better options.

But the GLP-1 field also raises massive payer questions. Obesity affects roughly 40% of American adults. If many of them take a drug costing $10,000-15,000 per year at list price, total spending would reach hundreds of billions annually. In a landmark move in August 2025, Novo Nordisk and Eli Lilly struck deals with the Trump administration to provide GLP-1 obesity drugs to Medicare at $245 per month -- the first time Medicare will cover anti-obesity medications. How payers respond more broadly will be one of the most consequential policy questions in healthcare over the next decade. The dynamic also illustrates the VC cycle perfectly -- in 2023-2024, every biotech VC wanted exposure to obesity and metabolic disease, and valuations for companies with GLP-1-adjacent programs spiked. The swarming behavior cascades through the entire capital structure.

## The drug that almost didn't happen

Keytruda (pembrolizumab) is one of the world's best-selling drugs, generating nearly $30B in 2024 for Merck. It's fundamentally changed cancer treatment and is approved across dozens of tumor types. But it nearly never made it to patients, and its backstory is a cautionary tale about how pharma M&A can nearly destroy transformative drugs.

The story begins in 2003 at Organon, a Dutch pharmaceutical company historically focused on women's health and neuroscience. Scientists were looking for antibodies that would stimulate PD-1 (a checkpoint receptor on

immune cells) to suppress the immune system, as a treatment for autoimmune diseases. Instead, they accidentally discovered PD-1 antagonists (antibodies that block PD-1), which unleashes the immune system to attack cancer cells. This accidental discovery is one of the great examples of serendipity in drug development.

The antibody was humanized in collaboration with MRC Technology (now LifeArc) starting in 2006. Then the corporate acquisitions began: Schering-Plough acquired Organon in 2007, and Merck acquired Schering-Plough in 2009 for $41B. At each acquisition, the PD-1 program was deprioritized. Organon's small team had seen it as high priority; Schering-Plough and then Merck didn't. By early 2010, Merck had essentially terminated the program and was preparing to out-license it. A term sheet reportedly valuing the program at essentially nothing was in place.

Then Bristol-Myers Squibb published promising clinical data on nivolumab (Opdivo), a competing PD-1 inhibitor, and Merck's leadership scrambled to re-evaluate what they were about to give away. The program was reactivated. Merck filed its IND by late 2010, received breakthrough therapy designation in 2013, and won FDA approval for advanced melanoma in 2014. Keytruda has since been approved for over 40 cancer indications and has become the backbone of modern cancer immunotherapy.

The lesson here is important. The M&A that drives biotech innovation can also destroy it when acquiring companies fail to recognize what they have. Keytruda survived because of persistent internal champions who kept arguing for the program, and because a lucky external signal from a competitor forced a re-evaluation. Many potentially transformative drugs don't get that lucky and are shelved, out-licensed for nothing, or simply forgotten in the transition from one corporate parent to another. This is one of the hidden costs of consolidation in the industry.

The Keytruda story also illustrates the molecule-target-pathway-effect framework. Pembrolizumab binds PD-1, a checkpoint receptor on T cells that normally acts as a brake on the immune system. Tumors exploit this by expressing PD-L1, which engages PD-1 and tells the T cells to stand down. Keytruda blocks this interaction, allowing the immune system to recognize and attack the cancer (like releasing a brake). It works across many cancer types because the PD-1/PD-L1 axis is a general immune evasion mechanism used by many tumors. This is why Keytruda has been approved for so many indications -- the target isn't cancer-specific, but immune-system-specific. Rather than designing drugs that attack specific cancer cells, checkpoint inhibitors attack the mechanism by which cancers hide from the immune system. It's a platform, not a point solution.

There's also an IP lesson embedded in the story. The three inventors (Gregory Carven, Hans van Eenennaam, and John Dulos) were recognized by the IPO Education Foundation as Inventors of the Year in 2016. Keytruda now has over 100 patent applications, the first of which are set to expire around 2028, setting up the next major patent cliff in the industry. Merck, like AbbVie before it, faces the challenge of replacing a $25B+ annual revenue stream.

## Closing

The biotech ecosystem is messy, expensive, and slow. Clinical trials fail more often than they succeed. Patents get weaponized to extend monopolies. Drug prices in the US are hard to justify to patients who can't afford their medications. Corporate bureaucracies nearly killed what became a blockbuster cancer drug. Promising science often dies between academic discovery and clinical development…

And yet, the system works. It works because it aligns the self-interest of scientists, entrepreneurs, investors, corporations, regulators, payers, doctors, and ultimately patients around a single, enormously difficult objective -- turning molecular insights into treatments that improve human life. The path from Gregory Winter's antibody phage display work in a Cambridge lab to a patient's cleared psoriasis, or from an accidental discovery at a Dutch pharmaceutical company to the most effective cancer treatment in history, is absurdly complex. But the fact that these paths exist at all, and that the system reliably produces new medicines every year, is one of the remarkable achievements of modern civilization.

The system isn't perfect, and the debates about pricing, access, patent strategy, and innovation incentives are real and important. But I think understanding the system is a prerequisite for improving it. You can't fix what you don't understand. And when you trace the full journey of a drug (from a university bench to a patient's medicine cabinet), what you see isn't a broken system, but a beautiful one that's working hard to solve problems that are very difficult. Importantly, it continues to evolve, with AI reshaping early discovery (and perhaps clinical success rates), gene therapies offering one-time cures, and the globalization of innovation expanding who can contribute, to name a few.

What strikes me most is the positive-sum nature of the whole thing. In the case of Humira: a scientist in Cambridge develops a technique and eventually wins a Nobel Prize; a biotech company is created; a pharmaceutical company acquires it and turns it into a global product; investors make returns; founders build organizations; doctors gain a new treatment option; patients with debilitating autoimmune diseases get their lives back. And eventually, when the patents expire, that molecule enters the public domain. Biosimilar manufacturers produce it, which causes prices to fall and access to expand. What was once proprietary becomes a permanent part of the world's medicine cabinet.

That's the machine: from a scientific lab to a public good, with every stakeholder in between getting rewarded for the value they added. And for all its friction and imperfection, no one has built a better one.