Breaking from the Herd: How Biopharma Can Escape Moloch's Trap

Until a few months ago, I hadn't given much thought to Moloch, the ancient figure associated in historical texts with extreme and costly sacrifices. But as I prepared for a conversation at SXSW with World Series of Poker champion turned game-theory evangelist Liv Boeree, I found myself going down a Moloch rabbit hole.

In her TED Talk and Win-Win Podcast, Boeree talks about something called "Moloch's trap," a concept that originates from Scott Alexander's seminal 2014 essay "Meditations on Moloch." In ancient accounts of Moloch, people are described as offering up things of immense value – most starkly, their own children! – not because such acts were desirable, but because they believed *making such sacrifices would avert something far worse*. Alexander uses Moloch as a metaphor for coordination failures: situations where individually rational, self-interested decisions by multiple actors collectively produce outcomes that are bad for everyone.

Boeree extends this metaphor further: *Moloch's trap is the name we give systems where no individual is the villain, yet everyone is complicit – and the consequences are severe*. It is, importantly, distinct from simple greed or bad intent. It is a *structural trap*, where even well-intentioned, intelligent actors get pulled toward destructive outcomes because the incentive architecture leaves them no choice. The structure itself makes detrimental choices feel rational. Ultimately, however, the structural trap creates a race to the bottom.

Preparing for the Fast Company Grill at SXSW sharpened my thinking about where this dynamic is most acutely playing out in the biopharma industry. It was Boeree who first prompted me to ask: does biopharma have a Moloch's trap? And if so, are consequences of that structural trap more dire than currently appreciated?

How might Moloch's trap apply to biopharma?

First, a quick primer on the five criteria that must be met to satisfy Moloch's trap:

1. **Escalating competition:** Are multiple parties locked in a race, where falling behind has severe consequences?
2. **Rational individual defection**: Does each actor have a compelling local reason to pursue the particular behavior?
3. **Collective self-destruction**: If every actor follows their local incentive, does the aggregate outcome harm everyone?
4. **Coordination failure**: Is there no credible, enforceable mechanism for all parties to agree to stop?
5. **Race to the bottom**: Does the competition drive quality, safety, or long-term value downward over time even as effort increases?

Now, let's apply these criteria to the biopharma industry.

Escalating competition: The drug development arms race is relentless. With patent cliffs looming and investors demanding near-term pipeline catalysts, every company must run faster just to stay in place. The stakes have never been higher for speed and success rates in clinical trials.

Rational individual defection: Following clinically validated, proven biological targets is an individually rational choice to address escalating competition. It reduces scientific risk: a drug

against the target worked before, so it will likely work again. And it increases speed: a playbook exists for how to do clinical development. Further, clinically validated targets offer cleaner regulatory pathways and satisfy investors who want de-risked assets. Every portfolio committee in pharma is incentivized to say "yes" to the validated target with scientific predictability and "no" to the moonshot that carries substantial risk.

Collective self-destruction: But when everyone follows the same clinically validated target, the aggregate result is duplicated R&D spend with near-zero marginal scientific value. This has the potential to cause market saturation that limits return on investment across the field. Moreover, there is massive opportunity cost: the novel targets that could cure Alzheimer's, autoimmune diseases, or treatment-resistant cancers go un-pursued. Worst of all, over time, there is risk of hollowing out of the industry's long-term innovation engine. At the extreme, this leads to more expensive R&D and less innovation – and an industry that collapses on itself.

Coordination failure: There is no industry-wide mechanism to say, "enough companies are already working on TROP2 in cancer – you go somewhere else." Intellectual property law, competitive secrecy, and antitrust concerns make coordination nearly impossible. Even if two heads of R&D privately agreed over dinner that target duplication was wasteful, their boards, investors, and legal teams would, quite rightly, prevent formal coordination. As you will see below, there are market forces that provide a mechanism to "coordinate," but these are implemented late in the journey of a medicine.

Race to the bottom: Ironically, as more companies pile onto the same target, each individual company must spend more to differentiate – faster trials, bigger clinical programs, more aggressive pricing strategies. The collective effort escalates while collective value creation stagnates.

Put together, all of this results in an ongoing problem for the biopharma industry: **target crowding**.

The risk of following the pack

Target crowding (or herding) is a canonical Moloch's trap. The data are stark, as I described recently in my Bull vs. Bear 2026 outlook post. A 2025 analysis by L.E.K. Consulting – a recent and comprehensive accounting of this problem – found:

- **Target crowding is increasing**
  - About a quarter of the approximately 13,600 drug-target pairs in the current preclinical and clinical R&D pipeline are concentrated around just 38 unique biological targets
  - Said another way, about 2% of active R&D targets – 38 targets in total – are associated with 50 or more drugs, representing 25% of ongoing pre-clinical and clinical R&D
- **Novelty is decreasing**
  - Novel targets entering the pipeline have declined sharply – from ~100 per year a decade ago to just 30 in 2024, even as the total R&D pipeline nearly doubled and Series A VC investment grew at ~18% annually
  - Between 2020–2024, only ~350 novel targets emerged across the entire industry

These data are consistent with an earlier (2020) analysis published in Nature, which found that 68% of targets pursued by top-10 pharmaceutical firms were the basis for 5 or more separate R&D programs. Others have highlighted the phenomenon of target crowding, too – for example, this 2025 blog from Bruce Booth, which builds on the L.E.K. study.

But how severe is the risk to the biopharma industry? Two examples highlight the magnitude of the problem.

The history of PD-1/PD-L1 therapies serves as a textbook example of target crowding. The checkpoint inhibitor revolution was a genuine breakthrough in cancer treatment. But what followed was a perfect Moloch cascade: dozens of additional PD-1/PD-L1 programs launched globally. Companies spent billions in clinical trials comparing their molecule to the same competitors, in the same patients, on the same endpoints. Many of these trials produced statistically similar but commercially undifferentiated results. Meanwhile, far less money flowed to other oncology mechanisms that might have helped cancer patients who don't respond to checkpoint inhibition – more than half, by some estimates.

A more recent example is the incretin class in obesity. Following the spectacular clinical validation of GLP-1 receptor agonism by semaglutide and tirzepatide, the pipeline response has been entirely predictable: dozens of companies are now chasing GLP-1, GIP, GCG, and combination incretin targets, each hoping that a marginal efficacy advantage will be enough to win. But this case adds a critical new dimension to the Moloch dynamic: the role of the market leader in structurally raising the barriers to entry.

Companies in the lead are not simply competing on science; these companies are compounding their advantage through manufacturing scale, supply chain and device integration, an expanding breadth of clinical programs, a rapidly accumulating data advantage, and novel direct-to-patient commercial channels, all of which are deeply capital-intensive and self-reinforcing. The result is that even a genuine "me-better" on efficacy may no longer be sufficient to enter and win in the real-world. The incretin race is not just an exercise in target crowding, but a masterclass in how a market leader can reshape the terrain so thoroughly that *Moloch's trap becomes, for most participants, inescapable*.

At its extreme, target crowding results in what I call pharmaceutical “involution,” a deliberate wordplay to contrast with pharmaceutical “innovation.” The physiological definition of involution is “the shrinkage of an organ in old age or when inactive.” In sociology, involution is used to describe exhausting competition that produces diminishing returns without meaningful forward progress – a society that expends ever-increasing effort to sustain a treadmill disguised as progress. In the real-world, pharmaceutical involution looks like dozens of PD-1/PD-L1 or incretin programs, billions spent, and a world where more than half of cancer patients still don't respond to checkpoint inhibitors, and where most patients with rare diseases have no approved therapy at all.

Amplifiers of Moloch’s trap

Target crowding and pharmaceutical involution do not arise in a vacuum. There are three forces today that materially amplify the Moloch dynamic in biopharma: capital markets, China, and investment in basic research.

Capital markets are not merely background noise. They are one of Moloch's primary engines. Capital preferentially flows into clinically validated targets because they offer legibility, reduced scientific risk, and faster paths to exit. Successful exits then attract more capital into the same spaces, and the cycle repeats. This feedback loop does not just reflect target crowding; it actively accelerates it.

China adds further fuel. Most assets emerging from Chinese biotech pipelines are directed at the same small set of clinically validated global targets, effectively multiplying the number of

competitors in already-crowded spaces. Much has been written on this dynamic elsewhere; the key point here is that it is a structural amplifier, not an anomaly.

Erosion of basic science funding is a third amplifying factor. In the US, the NIH – historically the world's largest public funder of biomedical research – has faced dramatic changes that put the distribution of funds at risk. Academic labs, which generate most truly novel target hypotheses, are already scaling back early-stage discovery work. This is directly relevant to the data cited above: the decline in novel targets entering the pipeline – from ~100 per year a decade ago to just 30 in 2024 – will presumably deepen as the upstream basic research engine is starved of fuel. Less basic science today means fewer novel targets tomorrow; it is a structural amplifier with a decade-long lag before its full consequences are felt.

From pharmaceutical involution to innovation

Moloch's trap is not a law of nature. It is a structural failure, and structural failures can be fixed.

Once Moloch's trap is recognized, there are escape routes. As Boeree describes, these involve (1) redesigned incentive structures, (2) building win-win rather than win-lose competitive frameworks, and (3) better coordination mechanisms. Of these three, I believe the second – creating win-win frameworks – is the most relevant for biopharma's target crowding problem.

When it comes to Boeree's first escape route, regulatory and market forces do, in theory, provide beneficial incentives in biopharma: commercial markets reward first-in-class (FIC) and best-in-class (BIC) medicines, and payers increasingly demand genuine differentiation. But these forces act far too late in the medicine development lifecycle – more than 10 years and billions of dollars after the crowding decision was made. If they were sufficient on their own, target crowding wouldn't exist. So, what else can we do?

When I first started down the Moloch's trap rabbit hole, I thought: "*Target crowding is the result of competitive intensity – no big deal. Market forces will correct it*." However, as I reflected further, I realized the consequences could be more severe and lead to pharmaceutical involution, which requires serious attention. The data make clear that market forces are not correcting this – they are accelerating it. That brings me to Boeree's second escape route: creating "win-win" frameworks to minimize the problem. Indeed, Boeree indicated this is one of her favorite solutions, as her aptly titled podcast, "Win-Win," suggests.

Here are specific actions we are taking at BMS to create win-win solutions and avoid Moloch's trap and the long-term consequences of pharmaceutical involution:

1. **Name the problem.** At BMS, we start by acknowledging target crowding as a Moloch's trap. When we do, we move from solely defining risk as scientific or clinical uncertainty. We expand risk to also include crowding risk, commercial differentiation risk, and long-term innovation risk. A novel target with high scientific uncertainty may carry less portfolio risk than the fourth entrant in a crowded space.

2. **Institutionalize "target diversity" as a key portfolio metric.** We treat the distribution of biological targets across our pipeline as seriously as other factors such as therapeutic area and modality diversity, clinical development cycle times, and portfolio efficiency. For example, we ensure that a meaningful and monitored share of our pipeline must be FIC. The exact percentage varies by therapeutic area, but it is generally 50% or higher. If BIC is our goal, then molecular differentiation must be substantial, and there must be a clear path to

test differentiation early in clinical development. In short: we make Moloch's intrusion visible on a dashboard.

3. **Adopt "causal human biology" as the north star – and defend it.** Our end-to-end R&D principles at BMS guide our portfolio decisions. The first principle – causal human biology – is the starting point for any new target. While clinical validation in humans is the strongest form of causal human biology, it also exacerbates pharmaceutical involution. Thus, we ensure that our portfolio has other sources of causal human biology (e.g., human genetics, longitudinal patient profiling) that offer the potential to be FIC, and we incorporate AI-driven target discovery to enable breaking from the herd. By surfacing novel biological targets that lack existing validation, our approach expands the frontier of therapeutic possibilities with options human researchers might otherwise overlook. This is a genuine Moloch escape route: biology-first selection naturally diversifies targets because diseases have many causal mechanisms. The key is to protect this principle from the gravitational pull of competitive crowding.

4. **Invest in therapeutic modalities to unlock novel and underexplored targets.** Our second R&D principle – matching modality to mechanism – provides an escape route by unlocking targets with strong causal human biology. As referenced in the L.E.K. study, over half of the ~5,000 "druggable" proteins remain untouched by drug development, and the remaining 15,000 "undruggable" proteins are not explored at all. BMS invests in novel therapeutic platforms like targeted protein degradation, cell therapy, and radiopharmaceuticals that have the potential to unlock targets with strong causal human biology.

5. **Have the discipline to stop programs early.** Our third R&D principle – path to clinical proof-of-concept (PoC) – creates a mechanism for pressure testing FIC and BIC designations assigned in the pre-clinical phase of a program. Avoiding Moloch's trap is not just about which programs you *start*. It is equally about when you *stop*. Too often in biopharma, programs are killed late, after hundreds of millions of dollars have been deployed, driven by sunk-cost reasoning, internal momentum, and the hope that the competitive landscape will thin itself out. This is deeply inefficient: capital trapped in late-stage losers is capital unavailable for genuinely innovative bets. At BMS, we have implemented a framework for earlier, more disciplined program termination – one that treats the decision to stop as a first-class strategic act, not a failure. As Boeree and I discussed at SXSW, a skilled poker player folds early when the odds turn, preserving capital for hands worth playing. The same logic applies to drug development. (For more on this, you can read my recent post on the Critical Value Creation Period.)

Boeree's final escape route is coordination. In pharma, this is not easy due to the competitive and proprietary nature of our business. However, possible coordination solutions include pre-competitive consortia (e.g., NIH and academic partnerships) that share target validation data before competitive programs launch; disease-focused public-private partnerships that map the target landscape openly; and advocacy for regulatory incentives (e.g., extended exclusivity, priority review vouchers) explicitly tied to target novelty. This final escape route also short-circuits one of Moloch's amplifiers: decreased funding in basic research.

Win-win for biopharma companies – and, most importantly, patients

Avoiding Moloch's trap in biopharma can unlock genuine scientific differentiation, delivering more FIC and BIC medicines to patients and providing novel cures for diseases that currently have no

treatments. The result avoids pharmaceutical involution and is a win-win for biopharma companies and patients alike.

Fundamentally, patients stand to benefit the most from a win-win solution because Moloch's trap has a profound human cost. When innovation systematically skews toward large, commercially attractive markets, patients with rarer diseases, more complex biology, or no response to existing therapies are left behind: the cancer patient who doesn't respond to checkpoint inhibition; the child with a rare metabolic disorder for whom the pipeline is effectively empty; the patient in a later line of therapy for whom the industry has moved on. These are not outliers. They represent the majority of unmet medical need. **Escaping Moloch's trap is, at its core, a patient imperative**.

As Boeree and I discussed, nothing is risk-free. Avoiding Moloch's trap through the win-win solutions above comes with its own potential for negative consequences. These include investor skepticism, first-mover risk, internal resistance, and competitive disadvantage in markets where a crowded target is the standard of care. There is also the lone defector problem – a dynamic from game theory – where if you escape Moloch but no other company follows, you may sacrifice short-term returns without changing the systemic dynamic. This is a genuine strategic risk.

This is where true leadership comes into play. As an R&D leader, I believe the strongest move is not just to resist Moloch within my organization. It's to use my voice to recruit other actors into a different game – a game where more players, and more importantly, more patients, win.

The choice is ours. We can continue to underestimate the risk of target crowding and pursue the same validated targets. Or we can elevate the risk to a Moloch’s trap and invest in courage and coordination to explore new and uncharted biological terrain. I have made my choice. The patients waiting for treatments that don't yet exist are counting on us to choose the latter.