Illustration credit: Soohyun Choi x Sandy Brian Hager

Canary in a K-hole

Ketamine is the blueprint for psychedelic treatment. It should be a warning.

By Sandy Hager

Jan 27, 2026

 **Ghost in the Megamachine**

https://sbhager.substack.com/p/canary-in-a-k-hole

Psychopharmacology has long been mired in a crisis of innovation. But in the past decade, new life has been breathed into the sector, and it has come from an unlikely source: ketamine.

Like other psychoactive drugs, ketamine has a colorful history. Originally developed as a general anesthetic for human and veterinary surgery in the 1960s, the drug became a staple in the 1990s rave scene.

By the late 1990s, a team at Yale University began researching ketamine as an agent that might mimic schizophrenic psychosis. But unexpectedly, they observed that participants, especially those with depression, experienced improvement in mood. This paved the way for a wave of findings over the next two decades that gave weight to the Yale team's accidental discovery of ketamine as a treatment for depressive symptoms.

In many ways, ketamine's arrival as a mental health treatment could not have been better timed. Millions of people around the world suffer from clinical depression. Up to a third of them do not respond to standard antidepressants. And when antidepressants do work, weeks can pass before they take effect, leaving those suffering in a dangerous limbo.

Against this backdrop, ketamine was celebrated in the media as a "wonder drug", the "anti antidepressant", and the most important advance in depression research in a half century.

These clinical breakthroughs have paved the way for a ketamine gold rush, with for-profit companies eager to cash in on the mounting hype. The number of clinics administering ketamine for "off label" use in the US has increased from around 60 in 2015 to between 500 and 750 in 2024. Changes in COVID-era legislation that made it easier to prescribe and distribute ketamine for at-home use via telehealth have fuelled the rise of mail-order ketamine companies.

The drug has also attracted growing interest from the pharmaceutical sector, with large established firms and biotech startups alike rushing to develop proprietary ketamine-based formulations. The first to receive regulatory approval for depression was Spravato (esketamine), developed by Janssen, a subsidiary of the pharmaceutical giant Johnson & Johnson (J&J). Approved by the US Food and Drug Administration (FDA) in 2019, Spravato has since grown into a blockbuster, the first drug in psychopharmacology to surpass $1 billion in annual sales since the launch of Selective Serotonin Reuptake Inhibitors (SSRIs) like Prozac in the 1990s.

With growing legal and illegal use, ketamine has come to occupy a distinctive place in contemporary culture. From Silicon Valley elites to the wider public, its dissociative effects have been taken up in a double role: as both a tool for productivity and creativity, and a temporary escape from the sensory overload and social isolation of digital life.

Just as ketamine has reenergised mainstream psychopharmacology, it has also provided a template for the broader psychedelic renaissance now unfolding in mental health care. Even though ketamine is an atypical psychedelic, pharmacologically and experientially distinct from classic psychedelics like psilocybin, LSD, or DMT, its clinical credibility and scalable treatment model have become the blueprint the rest of the field is now following. Listen to almost any psychedelic CEO and you'll hear the same refrain: ketamine is the lodestar; an almost obsessive point of reference for a sector that's faced a rocky path to commercialization.

Although ketamine has become a source of excitement for psychedelic CEOs, it should be far more unsettling for anyone concerned with patient outcomes. That's because, for all the early hype about ketamine as a wonder drug for depression, it is increasingly understood as a limited, short-term intervention, most useful in acute crises like suicidality but often requiring ongoing, potentially indefinite dosing to sustain its effects.

What's more, ketamine leaves patients squeezed between two bad options. Spravato is a patented, monopolised treatment that can still be punishingly expensive even with insurance and may be less effective than generic ketamine. The generic market is cheaper, but it is rarely covered by

insurance. And it involves off-label clinics with few standards or telehealth companies mailing powerful psychoactive drugs to people in vulnerable mental states with minimal screening or follow-up.

As ketamine use has become more mainstream, it has also become harder to ignore its darker side. What looked, at first, like a sleek new frontier in mental health care is increasingly shadowed by controversy: ketamine can be addictive, heavy use can cause severe bladder damage, and overdose is a real risk. Recent headlines, including those surrounding Matthew Perry's death, have made that painfully clear.

Rather than an inspiration, the commercialization of ketamine as a mental health treatment should instead be seen as a canary in the coal mine for the next wave of psychedelic medicine. If this is what success looks like, monopoly pricing on one end and a clinical free-for-all on the other, then psilocybin, MDMA, and the rest are unlikely to be any better for patients.

It's tempting to blame capitalism for ketamine's failure to live up to its early promise as a transformative treatment for depression. And no doubt capitalist markets and the profit motive are central drivers of both the hype and the harms. But the more interesting analytical question is why two radically different market structures, monopoly in drug development and cutthroat competition in clinical delivery, have nonetheless converged on the same disappointing outcome for patients.

A power-centred approach helps solve this puzzle. The problem is not simply that ketamine has been commercialised, but that commercialisation in capitalist healthcare is organised around what Thorstein Veblen called strategic sabotage: the deliberate restriction and channelling of productive capacity toward pecuniary ends.

In ketamine drug development, sabotage is most visible because power is concentrated: firms can use patents and regulatory exclusivities to restrict access, limit competition, and translate that control into differential earnings. In clinical delivery, sabotage takes a different form because power is diffuse and margins are thin. Providers compete by cutting the labour- and time-intensive elements of care such as screening, monitoring, integration, and follow-up, while expanding volume through marketing and convenience. These are very different market structures, but both subordinate therapeutic value to capitalization, and both therefore generate poor outcomes for patients.

To see how this works in practice, we need to pose a simple question: how, exactly, do the different players make money from ketamine as a mental health treatment?

Drug Developers: Control Versus Cure

Let's start with drug development. One route to profitability is to develop some new proprietary version of ketamine and ask regulators to approve it for specific psychiatric disorders. For pharmaceutical companies, the profit equation looks like this:

profit = (price per dose – production and R&D costs) × doses sold

But securing profits in this business is difficult. On average it takes 10 to 15 years for a drug to reach approval, at a cost of roughly $1-2 billion. And around 90 percent of drug candidates never make it to market.

To make matters worse, once a drug does get approved, the cost of producing each additional dose is usually very low, often close to zero. That makes pharmaceuticals easy to scale, but it also creates a vulnerability: in a competitive market, prices tend to get pushed down toward marginal cost. If rivals can make the same product for pennies, the price per dose collapses. And if the price collapses, the original investment becomes impossible to recover.

The aim of drug developers is therefore not productivity, but power. If a company cannot secure control over access and pricing, it cannot stabilize the "price per dose" term in the profit equation, and the entire investment case collapses. In that sense, the question investors care about is not only "will this molecule work?" but "will the firm have the exclusionary power needed to secure profits?"

Patents, regulatory exclusivity, and related protections provide that power. They are institutional mechanisms that restrict and channel productive capacity toward pecuniary ends. Instead of competing by making treatment cheaper and more widely available, firms compete by controlling the conditions under which treatment can be provided. Intellectual property is therefore not just a legal tool. It is the means by which a drug becomes a capitalized asset: a protected claim on future earnings.

Yet with ketamine the drive to secure IP is complicated, because the generic version is already available off label as a mental health treatment. The original patent expired in the 2000s, and the drug is now dirt cheap. Clinicians with a medical license can obtain it for one or two dollars a dose. So drug developers can't patent ketamine itself. They have to come up with a novel, non-obvious variation on a molecule that already works, and then convince the patent office that it's more than a minor tweak.

These peculiarities of profit-making in drug development also push the market toward concentration. The firms most likely to win are large incumbents that can finance years of uncertainty, survive multiple failures, and still have the resources to scale after approval. And once the drug is approved, first-mover advantage quickly becomes self-reinforcing. Regulatory approval brings institutional legitimacy, legitimacy brings uptake, and intellectual property locks in exclusivity for decades.

This concentration of power is exactly what we see in ketamine drug development. So far, the only ketamine-derived treatment approved for mental health is Johnson & Johnson's Spravato. To patent it, J&J isolated one of ketamine's isomers, esketamine, and developed it as a nasal spray, a more convenient method of administration than IV infusion. It is those seemingly minor tweaks combined with regulatory exclusivity that have created a blockbuster.

At first glance, ketamine drug development might look like a crowded field. Right now, there are around 30 companies other than J&J that either have a ketamine drug in clinical trials or hold

ketamine-related intellectual property (14 public and 19 private). But as Table 1 makes clear, the size difference is so extreme that it's hard to treat this as meaningful competition.

	Market Cap	R&D	Cash	Profit	Revenue	Financing
J&J	375	17	25	14	89	n/a
Others	32	0.7	1.2	0.07	5	0.04
Ratio (J&J / Others)	12	24	21	200	18	625

Table 1 Concentration in Ketamine Drug Development

Note: J&J stands for Johnson & Johnson. The numbers in the first two rows are in USD (millions). Market capitalization data are from July 31, 2025. R&D, cash, profit and revenues data are from 2024. Others for these items are the 14 other publicly listed ketamine drug development companies. Financing for others is the total lifetime fundraising of 19 private ketamine drug development companies. The numerator in the financing ratio is J&J's cash holdings in 2024.

Source: CB Insights; Datastream; Eikon Refinitiv

J&J's market capitalization is 12 times larger than the market cap of the 14 other publicly listed ketamine drug developers combined. It spends 24 times more on R&D than those firms combined. Its profits are roughly 200 times larger, largely because most clinical-stage drug companies run at a loss. And when you compare cash reserves to startup financing, the gap becomes almost comical: J&J's cash holdings are 625 times larger than the total lifetime financing raised by the 19 private ketamine companies combined.

With numbers like these, the conclusion is hard to avoid: J&J is dominant, and everyone else is fighting over scraps. The smaller ketamine developers are tiny by comparison, and with fewer resources they are less likely to make it over the line and win regulatory approval. Even those that defy the odds and reach the market are unlikely to have the capital, infrastructure, and institutional leverage to scale fast enough to compete with Spravato. For many, the best-case scenario is getting bought out late-stage by a larger pharmaceutical company that can finance commercialization.

J&J's monopoly matters because it produces predictable harms for patients. Generic ketamine is cheap; Spravato is not. For the lowest dose, the annual cost of Spravato for treatment-resistant depression can run over $55,000. Because it is the only FDA-approved ketamine-derived treatment, it is covered by most insurance plans, and for patients with good coverage the out-of-pocket cost can be relatively low.

But access is uneven. Not everyone has adequate insurance, and even those who do can find themselves trapped in bureaucratic gauntlets. Some patients report having to try multiple antidepressants first, and at least one plan has required electroconvulsive therapy before it will reimburse Spravato.

High prices and restricted access aren't the whole story. The deeper perversity is that the costly, insurance-backed option may not even be the best one. As evidence accumulates, Spravato increasingly appears less effective than generic ketamine in treating depression. That perversity is hard to make sense of if you assume drug markets reward the best therapies. But it makes sense in terms of strategic sabotage: the point is control, not cure.

Clinics: Cutting Care

Drug development is only one route to profit. The other is clinical delivery. Anyone with a medical license in the United States can open a clinic and prescribe generic ketamine off label for depression, anxiety, PTSD, addiction, or pain. This is how most ketamine therapy is administered.

For clinics, the profit equation looks like this:

profit = (reimbursement per session – operational costs) × number of sessions

The cost structure of running a ketamine clinic differs from drug development. Instead of enormous fixed costs and near-zero marginal costs, clinics face moderate fixed costs and relatively high marginal costs. Starting a clinic is expensive. You need real estate, equipment, licensing, and staff. But it's not the billion-dollar gamble of taking a molecule through clinical trials. Still every additional session requires staff time, room time, supplies, monitoring, and admin. That means growth is costly. Scaling up means adding more rooms, more clinicians, or more locations, and each step brings new expenses rather than effortless economies of scale.

Barriers to entry are also much lower. It is far easier to open a ketamine clinic than it is to run clinical trials and secure FDA approval for a new drug. And pricing power is weaker. Instead of patents that lock in exclusivity, clinics rely on harder-to-defend intangible assets to differentiate their products: brand, reputation, and marketing. These assets can provide advantages, but they are much harder to defend than intellectual property.

This profit strategy gives rise to a distinct market structure. If ketamine drug development tends toward monopoly, ketamine delivery tends toward fragmentation and competition. And that competition helps explain why ketamine clinics have struggled to turn a profit, especially when they try to scale.

During the psychedelic boom from 2019 to 2021, venture capital and private equity poured into ketamine clinic chains. But as the boom turned to bust, many discovered the constraint: clinics don't scale easily. Every new patient requires more staff time, more clinical space, and more overhead – all of which eats into the bottom line.

Table 2 captures the fallout. There have been fourteen companies of any notable size that have tried to build ketamine clinic chains. Of those fourteen, only four are still operating.

Company	Status
Awakn Life Sciences	Clinics closed in 2025 to focus on drug development
Braxia Scientific	Delisted, operations closed in 2025
Delic Holdings	Parent of KWC, operations closed in 2023
Emergence (Irwin Naturals)	Operations closed in 2024
Field Trip	Operations closed in 2023
Greenbrook TMS	Offers Spravato at TMS clinics
HOPE Therapeutics	Planning to acquire existing ketamine clinics
Ketamine Wellness Centers (KWC)	Operations closed in 2023
Novamind	Acquired by Numinus in 2023
Numinus	Acquired by Stella in 2024
Pasithea Therapeutics	Clinics closed in 2023 to focus on drug development
Revitalist	Cease to trade notice, downsized from 10 to three clinics
Stella	Operates clinic network, acquired Numinus in 2024
Wellbeing Digital Sciences	Cease to trade notice, operations closed in 2022

Table 2 Ketamine Clinics: Operational Status

Source: Author's compilation based on company filings, press releases, bankruptcy notices, acquisition announcements, and media reporting (2022–2025).

For patients, the picture is grimly familiar. One of the real advantages of generic ketamine is that it is both cheaper and often more effective than Spravato. An IV session of generic ketamine typically costs around $400 to $800. That's far less than Spravato, but it's still unaffordable for many people, especially because most insurance plans won't cover generic ketamine for depression. It is not FDA-approved for that indication. In practice, this means that patients can end up paying more out of pocket for generic ketamine than they would for Spravato, assuming their insurance will cover Spravato at all.

Then there is the problem of oversight. Because most ketamine clinics operate off label, there is no consistent regulatory framework governing who qualifies for treatment, what dosing protocols should look like, what monitoring is required, or what follow-up care patients should receive. The result is a fragmented "Wild West" market in which standards vary widely from provider to provider. Many clinics are run by anesthesiologists or pain specialists with little training in psychiatric care, even though they are treating patients with severe mood disorders.

And the financial pressures built into the clinic model push in predictable directions. To stay afloat, providers have to do two things: increase patient volume and contain operational costs. Clinics compete for customers through aggressive marketing, often presenting ketamine as a miracle cure while downplaying risks and uncertainties. At the same time, the labour- and time-intensive elements of treatment are the easiest to cut: thorough screening, careful monitoring, recovery time, and meaningful follow-up.

In practice, this can mean thin staffing, rushed appointments, and minimal aftercare. Some clinics monitor multiple patients at once, sometimes remotely via cameras. Others treat several patients in the same room. The setting may look professional, even luxurious, but the logic is closer to throughput medicine: keep the chairs filled, keep sessions moving, keep costs down.

What looks like bad practice is better understood as a business logic. In a fragmented clinic market, providers can't raise prices, so they cut costs. And in ketamine care, cost cutting often means cutting care: staff time, clinical attention, and continuity. In this context, strategic sabotage doesn't restrict access so much as degrade it. Treatment expands, but care gets hollowed out.

So even though the delivery market is competitive rather than monopolistic, the result can be the same: patient wellbeing gets subordinated to the demands of profitability.

The most extreme version of ketamine delivery is at-home treatment. During COVID, changes in US telehealth policy made it easier to prescribe controlled substances remotely. That opened the door to a wave of mail-order ketamine companies: platforms that recruit patients online and ship ketamine, usually oral lozenges, directly to their doorstep.

One of the biggest players is Mindbloom. If a patient is accepted, they receive the company's trademarked Bloombox, complete with eyeshades, a journal and pen, a blood pressure monitor, and ketamine lozenges. The pitch is simple: ketamine therapy without the clinic. Cheaper, more convenient, and packaged as a guided transformation rather than a medical intervention.

The profit equation here is the same as in-person delivery, but the cost structure is different. Telehealth companies don't have to pay for treatment rooms, recovery spaces, or in-person staffing. That makes the model far easier to scale. But barriers to entry are also even lower, competition is intense, and pricing power is weak. Without patents, mail-order ketamine companies rely on branding and wellness lifestyle aesthetics to differentiate what is, pharmacologically, a generic drug.

For patients, the main advantage of mail order ketamine is cost. Mindbloom charges around $159 per session. But because these services are rarely covered by insurance, even the "affordable" option remains out of reach for many. And the lower price comes with trade-offs. The entire model depends on widening the distance between patient and provider: ketamine is taken at home, often with limited supervision, outside the clinical setting.

The same pressures that shape clinics show up here, often amplified. Companies chase volume through slick advertising and glowing testimonials that blur the line between marketing and medical advice. In one Mindbloom testimonial, a patient claims that a few sessions were like "five years of therapy." Meanwhile, patients have reported cursory screening, sometimes a brief video call followed by a large supply of ketamine arriving at their door the next day.

Cost cutting also takes a familiar form: watered-down care. Telehealth patients describe minimal preparation, limited follow-up, and difficulty reaching a human when things go wrong. One New York Times account is particularly harrowing. A woman began taking ketamine through a mail-order company called Joyous, felt worse, and became suicidal. For ten straight weeks, she reported suicidal thoughts in her responses to a company questionnaire. Each time, she received the same automated text message in reply. It included the national suicide hotline and ended with: "We hope you start feeling better very soon!" She claimed that not once in those ten weeks did a human being check on her.

Most recently, Mindbloom has been swept up in controversy of its own. Until recently, mail-order ketamine companies focused on lozenges. But earlier this year, Mindbloom announced that it would begin offering at-home injectable ketamine. Injectable ketamine has a faster onset and more intense dissociative effects than oral dosing, and it carries greater risks, including blood pressure spikes and a higher potential for misuse. Critics argue that at-home injections cross a line: they are simply too dangerous, especially for people in fragile mental states.

Telehealth is the logical conclusion of the cost-cutting pressures built into ketamine care. Where brick-and-mortar clinics reduce care to survive, mail-order ketamine removes the clinic altogether. Supervision and follow-up cease to be standard features of treatment and become discretionary costs. This is the same strategic sabotage, now operating with fewer constraints and shifting greater risk onto patients.

The Canary

Ketamine's commercialization has revived a stagnant field and helped legitimize the wider psychedelic renaissance. It is the first serious test case for what "psychedelic medicine" looks like when it is scaled. And what it shows is not a miracle cure, but a system that often produces disappointment and harm.

These outcomes don't depend on any single market structure. In drug development, power is concentrated. Patents and regulatory exclusivities restrict access, limit competition, and turn a treatment into a protected earnings stream. In clinical delivery, power is diffuse. Thin margins and competition push providers to cut labour-intensive elements of care while expanding volume through branding, marketing, and convenience. Diverse markets converge on the same result: therapeutic value is subordinated to capitalization, and patients are left to absorb the fallout.

This is why ketamine is a canary in the coal mine for the next wave of psychedelic medicine. If ketamine is the blueprint, then psilocybin, MDMA and DMT are unlikely to escape the same structural pressures. We should expect monopoly dynamics wherever firms can secure intellectual property and regulatory privilege, and we should expect a clinical free-for-all wherever treatment is delivered off label or through loosely regulated service models. In both cases, "innovation" is defined by profitability, not care or wellbeing.