SOUTH BEACH LONGEVITY
Peptide Market & Culture Intelligence

SBL INTELLIGENCE REPORT · THERAPEUTIC PEPTIDES · AUGUST 2026

From the Lab to the Feed

How therapeutic peptides entered medicine, markets, and culture, and which molecules are moving through that pipeline now

Peptides are having a moment, but "peptides" is really several moments happening at once, and they are easy to confuse. This South Beach Longevity intelligence report separates the parts: a pharmaceutical revolution, a scientific field, a gray online marketplace, and a wellness subculture. Drawing on tens of thousands of news articles, social posts, live retail-market data, and the public science registries, it asks a single question of the evidence and answers it in depth for both the general reader and the professional investor: beneath the GLP-1 weight-loss boom that dominates the headlines, is a broader peptide movement genuinely taking hold, and where is it going?

COMPILED BY South Beach Longevity · August 2026
EVIDENCE ~32,700 news articles · ~18,100 social posts · a live research-use-only retail-price corpus · ClinicalTrials.gov and PubMed registries METHOD Descriptive analysis of a collected corpus, class-labelled by evidence type CONSTRAINT This is a market, media, and culture study. It recommends no dose, route, or human use of any compound in any person, and is not medical advice.

How to read this document. This report mixes very different kinds of evidence, and it is written to keep them visible rather than blur them together. When it says *people are reporting* something, that means a claim appeared in public discourse, a post, a comment, a news quote, and nothing more. When it says something has been *shown* or *demonstrated*, that points to a trial or a published result. **Search interest, social attention, media coverage, commercial availability, and clinical proof are five different things**, and a compound can rank high on one while scoring near zero on the others. Every chart names its source, its date range, and what the number does and does not mean. Where the data cannot answer a question, the report says so instead of guessing.

Executive summary

For most of their history, therapeutic peptides were a specialist concern, a class of drugs, a corner of biochemistry, a set of injections that endocrinologists and a few bodybuilders understood and almost nobody else thought about. In the space of a few years that changed completely. The reason is mostly two molecules: **semaglutide** and **tirzepatide**, the active ingredients in Ozempic, Wegovy, Mounjaro, and Zepbound. They turned "peptide" into a word that shows up in stock analysis, late-night monologues, and grocery-industry earnings calls. But the same wave lifted a second, stranger world into view: a loosely regulated online market in "research" peptides with names like BPC-157, TB-500, and GHK-Cu, sold for laboratory use but bought and injected by a growing wellness and fitness subculture.

This study draws on a large collected evidence base, roughly **32,700 news articles** and **18,100 social-media posts** tied to specific compounds, a live price corpus covering **302 U.S. research-use-only retailers**, a market-strength model, and public trial and publication registries, to map how peptides crossed from science into commerce and culture, and to test whether the boom is as broad as it looks. The central findings:

- **The news and the internet are telling two different stories.** Of the news articles the system could tie to a specific compound, **about 93% mention a GLP-1 drug**. Of the social-media posts it could tie to a compound, **about 80% name a *non*-GLP-1 peptide** and only ~8% name a GLP-1 drug. The mainstream press is almost entirely a GLP-1 story; the peptide *internet* is almost entirely about everything else.
- **There is a real non-GLP-1 peptide movement, but it lives on social media, not in the newspaper.** Compounds like GHK-Cu, BPC-157, TB-500, MOTS-c, and KPV draw hundreds of social posts each while attracting little or no mainstream coverage. The movement is real; it is just not where journalists have been looking.
- **Different peptides took opposite roads to prominence.** GLP-1 drugs diffused *top-down*: clinical trials, then the business press, then the public. Research peptides diffused *bottom-up*: a fitness and biohacking community talking on TikTok and Reddit years before any newsroom noticed. BPC-157 is the clearest crossover: online since at least 2019, but essentially absent from the news until mid-2024, when coverage finally began to chase the culture.
- **The commercial market ranks these worlds as equals; the science does not.** In the RUO retail market, "research" peptides such as MOTS-c, BPC-157, and GHK-Cu sit at the very top of a market-strength model, right beside the GLP-1 drugs, while carrying a fraction of their clinical evidence. The decoupling is stark at the extremes: BPC-157 is stocked by more online sellers than semaglutide, 5-Amino-1MQ carries hundreds of live listings on a base of just three lifetime publications and zero human trials, and orforglipron draws 2,786 news mentions yet has only a token gray-market presence. The market is pricing attention, not evidence.
- **For an investor, the signal is in the structure, not the size.** No audited market-size figure can be derived from this evidence, and the report refuses to invent one. But the shape is legible. The gray market is fragmented and low-moat: the same compound at the same dose varies severalfold in price across sellers, and its margin sits in vialing and branding rather than in manufacturing, since bulk powder trades per gram while retail vials sell per milligram. The durable pharmaceutical value sits with the incretin innovators and the regulated "picks and shovels"

around them (compounding pharmacies, telehealth platforms, testing labs), and the entire research-chemical layer rests on a "research use only" status that a single regulatory shift could redraw.

- **The most-discussed *cultural* effects of peptides are claims, not findings.** The corpus contains more than 8,000 tagged observations of GLP-1 drugs supposedly changing everything from drinking and addiction to grocery spending and airline fuel costs. Across every one of those categories, the share carrying established causal evidence is **zero**. The society-level story is, so far, a story people are telling, an important cultural fact in its own right, but not a demonstrated one.

The picture that emerges is not a single "peptide boom" but a set of overlapping ones moving at different speeds and on different evidence. One is a genuine pharmaceutical revolution with real trials behind it. Another is a fast, largely unproven consumer market running on word of mouth. Keeping them apart is the whole point of this report. The chapters that follow take each in turn: the science of the incretin drugs and the pipeline behind them; the eight classes of evidence, from peer-reviewed literature to TikTok; the people and subcultures driving adoption; a dedicated investor and professional lens on the value chain, competitive structure, and regulatory risk; a compound-by-compound dossier of sixteen molecules; and a closing read on where each of these worlds is heading.

PART I

The phenomenon

01 · The peptide moment

Type "peptide" into a search box today and the autocomplete tells the story: peptides for weight loss, peptides for skin, peptides for healing, peptides for muscle, "peptides before and after." A decade ago the same word would have returned biochemistry textbooks. Something moved a technical term out of the laboratory and into the vocabulary of ordinary consumers, and it moved fast.

A peptide is, in plain terms, a short chain of amino acids, the same building blocks that make up proteins, just fewer of them strung together. That biological fact matters less to this report than a social one: **the word "peptide" now points at several different worlds that outsiders routinely mistake for each other.** There are peptide *medicines* approved by regulators and prescribed by doctors, insulin being the century-old original. There are peptide *drugs still in development*, tested in humans but not yet approved. There are *compounded* versions mixed by pharmacies to fill shortages or personalize a dose. And there is a large, largely online market in peptides sold "for research use only", legal to sell as laboratory chemicals, not approved for people, but bought and used by a wellness and fitness subculture anyway.

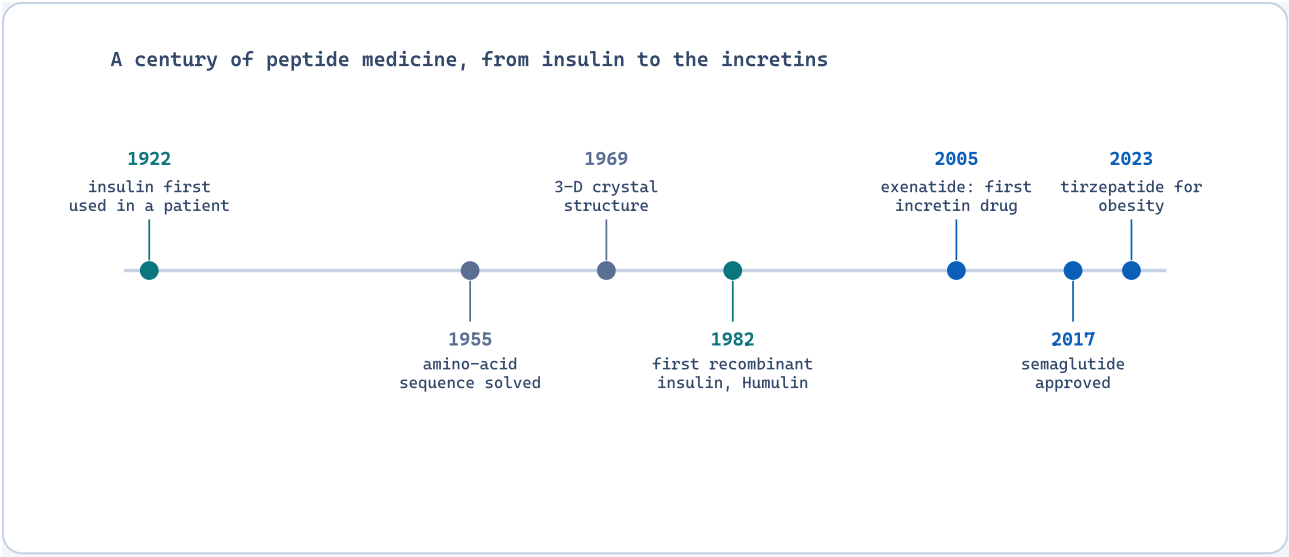

FIGURE 1 A hundred-year arc. The GLP-1 boom is the latest chapter in a century of peptide medicine that began with insulin in 1922 and ran through the first genetically engineered drug, recombinant insulin in 1982, to the incretins. Source: SBL insulin and incretin monographs.

These worlds share a word and very little else. A patient injecting FDA-approved semaglutide under medical supervision and a forum user reconstituting a vial of gray-market BPC-157 in their kitchen are both, technically, "using peptides." Treating them as the same phenomenon, or assuming that because one is rigorously studied the other must be too, is the single most common error in public conversation about this topic. Much of what follows is an attempt to hold the categories apart.

One force sits behind the whole moment and has to be named at the outset. The GLP-1 weight-loss drugs are so enormous, commercially, medically, and in sheer volume of attention, that they distort any attempt to see the rest of the field clearly. In the collected news corpus, GLP-1 drugs appear in roughly nine of every ten compound-linked articles. They are simultaneously *part* of the peptide story and large enough to eclipse it, the way a full moon washes out the surrounding stars. A recurring task of this report is to put the GLP-1 phenomenon behind a filter and ask what the sky looks like without it.

02 · What counts as a therapeutic peptide?

Before measuring a phenomenon it helps to define its edges, and here the edges are genuinely fuzzy. The taxonomy this study uses comes from the underlying research catalog, which tracks **301 distinct compounds** grouped by function rather than by chemistry, a practical scheme that reflects how these molecules are actually marketed and discussed. The largest groups are revealing in themselves:

FUNCTIONAL GROUP	COMPOUNDS TRACKED	EXAMPLE MEMBERS
Nootropic / neuro	28	Semax, Selank, cognitive research compounds
Metabolic / weight	25	Semaglutide, tirzepatide, retatrutide, 5-Amino-1MQ

FUNCTIONAL GROUP	COMPOUNDS TRACKED	EXAMPLE MEMBERS
SARMs and adjacent	21	Selective androgen receptor modulators (not peptides, but sold alongside)
Growth-hormone axis	20	Ipamorelin, CJC-1295, sermorelin, tesamorelin
"Khavinson" bioregulators	20	Epithalon and the Russian peptide-bioregulator family
Cosmetic	18	GHK-Cu and skin/hair peptides
Reproductive axis	18	Kisspeptin, PT-141, gonadotropin-related
Longevity / bioregulator	18	Pineal and organ-specific peptides
Muscle / growth, repair, recovery	19	BPC-157, TB-500, IGF-1

Note on categories. This is a market taxonomy, not a pharmacological one. It deliberately includes things a chemist would exclude, SARMs are not peptides, NAD+ is a nucleotide, because they are sold, stacked, and talked about inside the same subculture. The grouping reflects the world being studied, not a claim about molecular identity.

Two things stand out. First, only a small share of these compounds, the GLP-1 and closely related metabolic drugs, have anything like a mainstream medical profile. The overwhelming majority are what the underlying data flags as "research chemicals": legal to sell for laboratory use, not approved for humans, and known mainly to a specialist online audience. Second, the categories that *dominate the catalog* (nootropics, growth-hormone secretagogues, bioregulators, cosmetic peptides) are almost entirely absent from the categories that *dominate the news* (obesity and diabetes). The gap between what is sold and what is covered is a preview of this report's main finding.

It is worth stating plainly what "therapeutic peptide" will and will not mean here. This report treats the whole marketed universe as its subject, approved drug, investigational candidate, and gray-market research chemical alike, because the phenomenon under study is precisely the *blurring* of those lines in public understanding. But it labels which is which at every turn. When a compound has been through human trials, the report says so. When its reputation rests entirely on animal studies and self-reported anecdotes, the report says that too, in the same sentence.

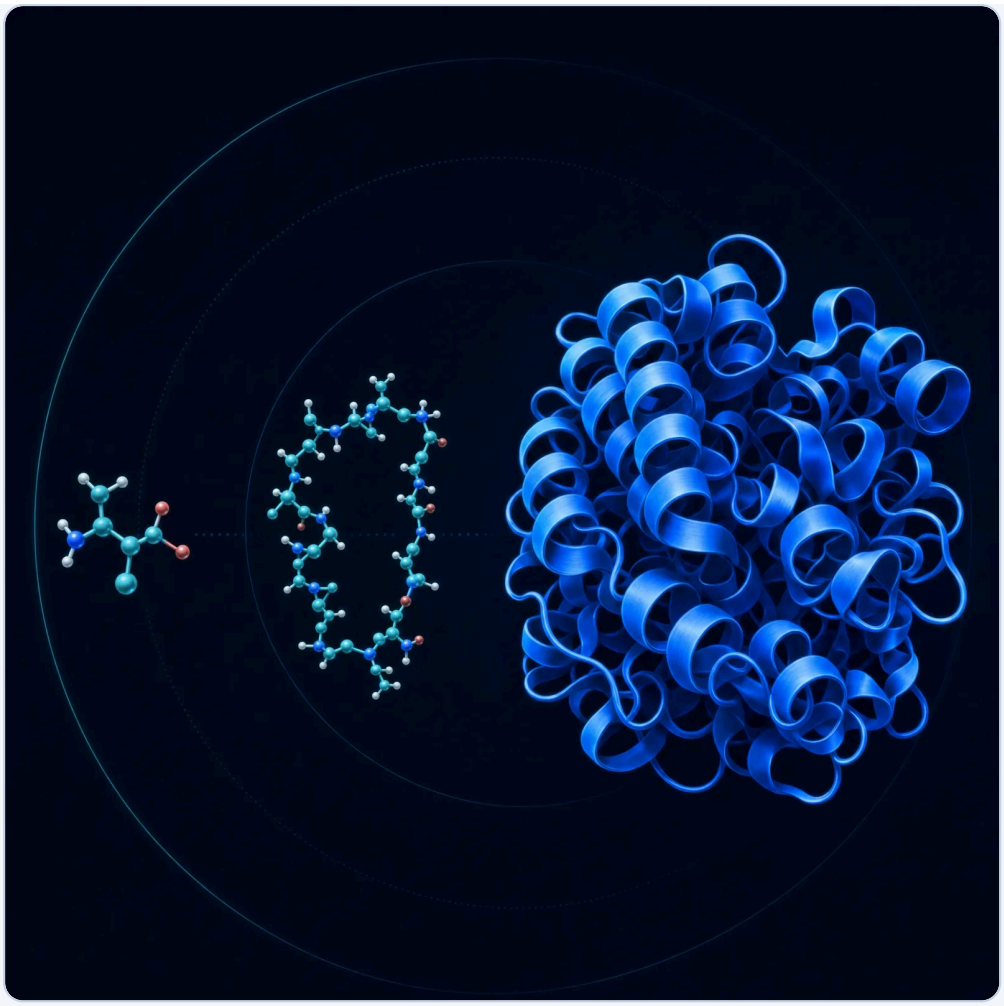

FIGURE 2 From amino acid to peptide to protein. A peptide is a short chain of amino acids; a protein is a much longer, folded chain of the same building blocks. The difference is size and folding, not a hard line. Illustrative scientific concept, not drawn to exact scale.

03 · How peptides left the lab

Ideas and products move from specialist worlds into mainstream life along fairly well-worn paths. A new molecule might travel: laboratory research → clinical trials → regulatory approval → medical prescribing → news coverage → search interest → social discussion → consumer purchase. That is the textbook route, and it is roughly the route the GLP-1 drugs took.

But it is not the only route, and the peptide story is interesting precisely because different compounds took different ones. A second path runs in nearly the opposite direction: a fitness or biohacking community starts experimenting → discussion spreads on social media → "research use only" vendors list the compound → search interest climbs → and only much later, if ever, does mainstream journalism or formal science take notice. The first path is *top-down*, pushed by pharmaceutical companies and regulators. The second is *bottom-up*, pulled by consumers and online communities ahead of the institutions.

This report organizes its evidence around the stages those paths pass through, treating each as a distinct signal that can be measured separately:

- **Science**, what is being researched and published.

- **Clinical development**, what is being tested in humans.
- **Medicine**, what is actually prescribed and used under supervision.
- **Commerce**, what is being sold, and through which channels.
- **Search**, what the public is trying to learn about.
- **Social culture**, what people are discussing, recommending, and reporting.
- **Journalism**, what has entered mainstream conversation.
- **Consumer adoption**, evidence that people are actually buying and using these products.

The key discipline is to never let one stage stand in for another. A compound with thousands of social posts and no clinical trials has *cultural* momentum, not *medical* validation. A drug with heavy news coverage but modest search interest may be a story journalists find important rather than one the public has embraced. The rest of this report walks the stages one at a time, and then, in the sections on individual compounds, puts them back together to show how a given molecule actually traveled. As will become clear, the shape of that journey is often the most informative thing about it.

PART II

The GLP-1 revolution

04 · The incretin machine, explained

To understand why two injections rewired an entire drug class, and a chunk of the stock market with it, you have to start inside the gut. After a meal, the small intestine releases a set of signaling molecules called *incretins*, hormones whose job is to tell the rest of the body that food has arrived. The most important of them is *GLP-1*, short for glucagon-like peptide-1. GLP-1 does several useful things at once: it prompts the pancreas to release insulin, but only when blood sugar is actually high, so it lowers glucose without the crash-risk of older diabetes drugs; it slows *gastric emptying*, the rate at which the stomach passes food onward, so you feel full longer; and it acts directly on appetite centers in the brain, turning down hunger and what users describe as "food noise." A second incretin, *GIP* (glucose-dependent insulinotropic polypeptide), works alongside it on insulin release and fat handling. A third hormone, *glucagon*, is in some ways GLP-1's opposite: it raises blood sugar and, at the right dose, lifts the body's overall energy expenditure, the rate at which it burns calories at rest.

The natural versions of these hormones are useless as drugs for one boring reason: the body destroys them in about two minutes. An enzyme called DPP-4 clips GLP-1 apart almost as fast as the gut makes it. The whole pharmaceutical achievement of the last two decades is *engineering around that clock*. Chemists rebuilt the GLP-1 molecule so the enzyme can no longer grab it, and attached fatty-acid chains that let it cling to albumin, a carrier protein in the blood, so it lingers for days rather than minutes. *Semaglutide*, the molecule in Ozempic and Wegovy, is dosed once a week. That is the difference between a hormone that flickers after lunch and a drug that holds appetite down around the clock. Sustained, unrelenting suppression of hunger, not any exotic new mechanism, is what produces weight loss on a scale earlier diet drugs never touched.

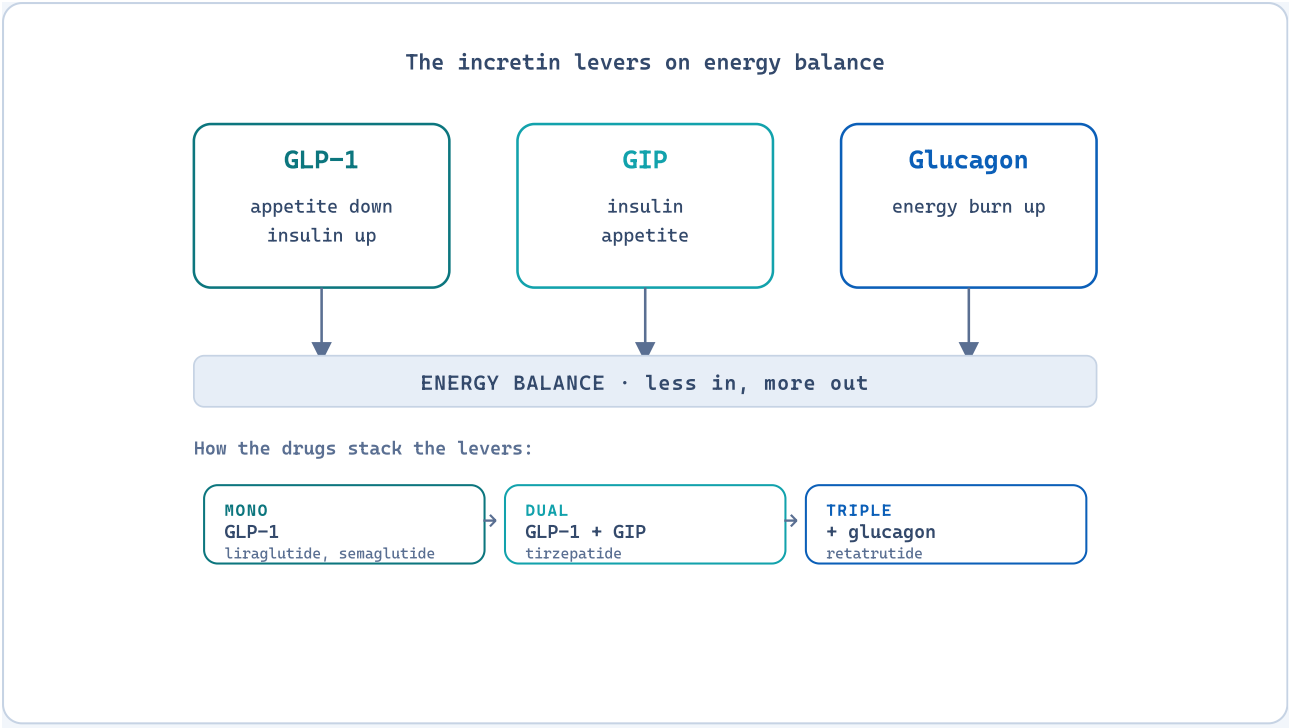

FIGURE 3 The incretin levers. GLP-1 and GIP mainly cut how much you eat and sharpen insulin release; glucagon raises how much energy you burn. The drugs escalate from pulling one lever (semaglutide) to two (tirzepatide) to three (retatrutide). Schematic, not to scale. Source: SBL incretin monographs.

This is why the newest molecules push further by adding targets rather than raising the dose. *Tirzepatide* (Mounjaro, Zepbound) is a *dual agonist*: one molecule that switches on both the GLP-1 and the GIP receptors. *Retatrutide*, still investigational, is a *triple agonist*, hitting GLP-1, GIP, and the glucagon receptor at once, the last of which is thought to raise calorie burn on top of cutting intake. The plain-language version is that the field has moved from pulling one lever to pulling two, then three, on the same wire of energy balance, and each added lever has, so far, deepened the effect. That escalation is the engine of the pipeline described next, and it is also the reason investors treat each new molecule as a potential successor rather than a copy.

WHY THE SCIENCE MATTERS COMMERCIALLY

The value here is not a clever receptor. It is durability at scale. A weekly injection that reliably reduces body weight addresses obesity and type-2 diabetes, two of the largest chronic conditions in medicine, with a compliance profile ordinary people can actually keep. Every element of the market story that follows, the duopoly, the shortage, the investor framing, flows from that single fact of biology: the effect is large, it is sustained, and the addressable population is enormous.

05 · The incretin pipeline, by the numbers

The incretin drugs are best read not as a handful of products but as a *pipeline*, a generational sequence in which each molecule is more ambitious than the last and the evidence base compounds behind it. Two public registries let us measure that pipeline directly. *PubMed* indexes the biomedical literature; *ClinicalTrials.gov* registers human studies. Queried compound by compound [PubMed], [ClinicalTrials.gov], they show a field that is not merely large but accelerating, and they show exactly where the acceleration is happening.

The headline totals are these. The nine GLP-1 and incretin compounds tracked here carry roughly **20,000 PubMed-indexed publications and about 2,196 registered human trials between them** [PubMed], [ClinicalTrials.gov]. Semaglutide alone anchors **741 trials and 5,465 publications**, of which more than **1,000 appeared in 2024** in that single year. Semaglutide and liraglutide together (**1,254 trials**) outnumber every non-GLP-1 research peptide in this study combined by roughly an order of magnitude. This is not a marketing phenomenon dressed as science; it is one of the most heavily investigated areas in all of modern pharmacology.

The generational shape is the more interesting finding. Line the family up oldest to newest:

COMPOUND	MECHANISM	STATUS	REGISTERED TRIALS	PUBMED (LIFETIME)	PAPERS IN 2024
Exenatide	GLP-1 (first-gen)	Approved	377	4,489	196
Liraglutide	GLP-1 (daily)	Approved	513	5,868	567
Dulaglutide	GLP-1 (weekly)	Approved	139	1,310	178
Semaglutide	GLP-1 (weekly)	Approved	741	5,465	1,052
Tirzepatide	GLP-1 / GIP (dual)	Approved	274	2,391	399
Cagrilintide	Amylin (CagriSema combo)	Investigational	43	97	14
Survodutide	GLP-1 / glucagon (dual)	Investigational	26	75	21
Orforglipron	Oral GLP-1 (small molecule)	Investigational	50	127	17
Retatrutide	GLP-1 / GIP / glucagon (triple)	Investigational	33	169	45

Source: ClinicalTrials.gov API v2 (registered-trial counts = development activity, not approval or usage) and NCBI PubMed (publication counts = research activity, not quality), queried by intervention/compound name, August 2026. Counts for older names can be inflated by synonym expansion; incretin names are relatively clean.

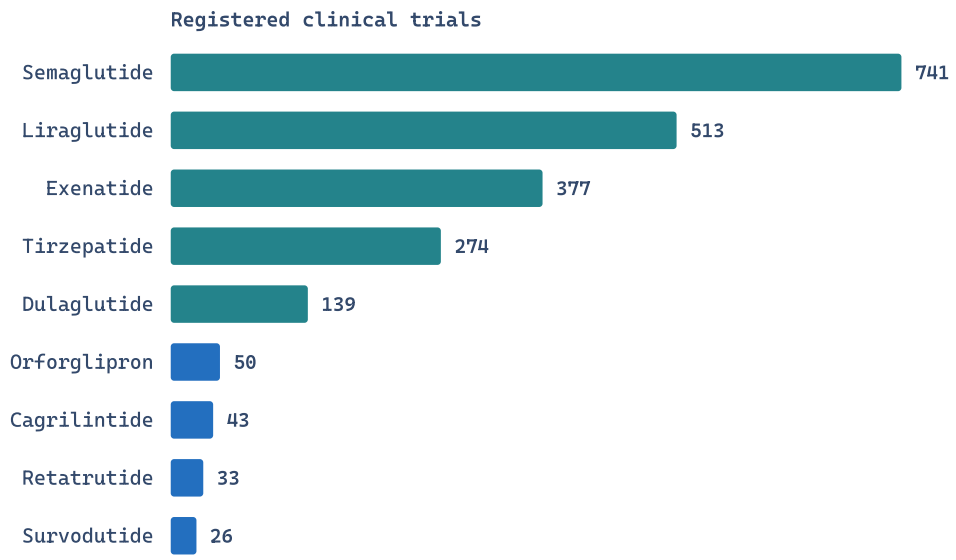

FIGURE 4 The incretin pipeline. Registered clinical trials for the GLP-1 and incretin family, teal for approved drugs and blue for investigational candidates. A deep, multi-generation development program sits behind the headline drugs. Source: ClinicalTrials.gov.

Read the columns together and the curve tells a story. *Exenatide*, the first-generation drug, has a large lifetime literature but only 196 papers in 2024: a mature molecule whose research wave has crested. *Semaglutide* and *tirzepatide* are at full boil, semaglutide's 1,052 papers in a single year is more than exenatide accumulated across several. And the newest entrants, retatrutide, survodutide, orforglipron, cagrilintide, still carry small lifetime totals (dozens to low hundreds) but a strikingly high share of their entire literature is from 2024, the signature of programs still on the steep part of the climb. GLP-1 research was near-flat before roughly 2018 and has risen almost vertically since [PubMed]. This is what a genuinely well-funded scientific boom looks like from the inside: not one drug, but a relay in which each molecule hands attention and capital to a more ambitious successor.

Two entries deserve a plain-language footnote because they define where the pipeline is heading. *Orforglipron* is an **oral** GLP-1 agonist, a pill rather than an injection, and a small chemical molecule rather than an engineered peptide. If it clears trials, it removes the needle, the single biggest friction point in the entire GLP-1 experience, and that is why it already draws heavy coverage (2,786 news mentions) despite not being approved for anyone [news corpus]. *CagriSema* pairs semaglutide with *cagrilintide*, an *amylin* analogue (amylin is yet another satiety hormone), a bet that combining hormone systems beats maximizing any one of them. The direction of travel is clear: more targets, easier delivery, deeper effect.

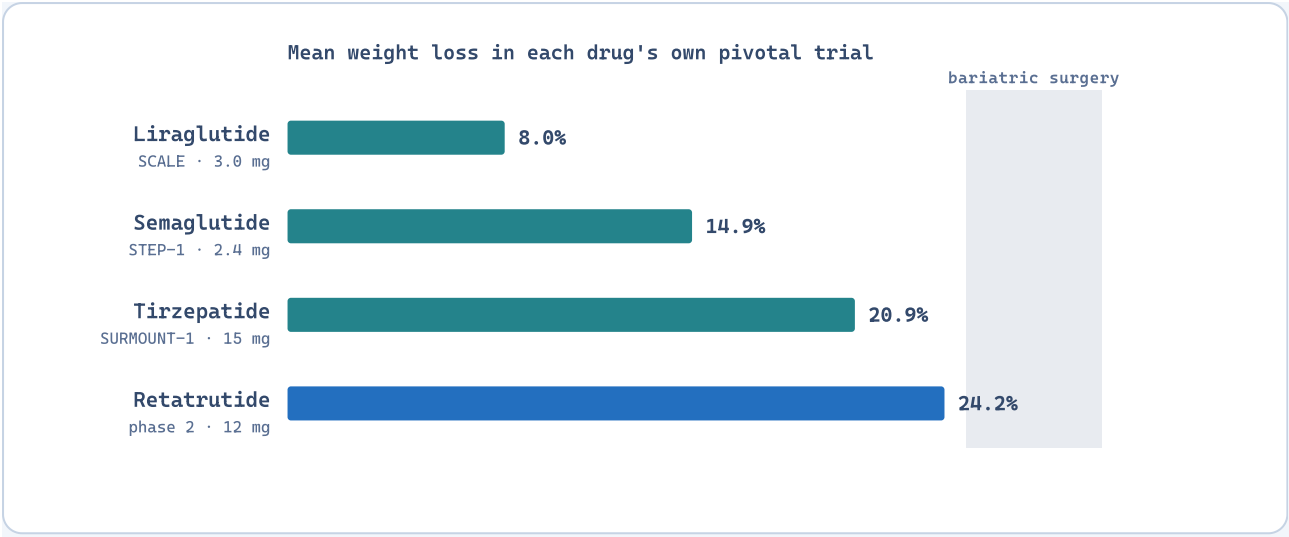

FIGURE 5 The generational weight-loss ladder. Mean body-weight reduction for each drug in its own pivotal trial, with the bariatric-surgery range shown for scale. These are separate trials, not a head-to-head comparison. Sources: SCALE, STEP-1, SURMOUNT-1, and the retatrutide phase-2 trial.

06 · Why this became an investor story before a medical one

Here is the fact that most cleanly separates the peptide story from an ordinary medical advance: for most of the public, it arrived as a *stock* story. When the collected news corpus is sorted by the outlet that published each article, the outlets covering peptides most heavily are not health desks. They are financial-news operations. **CNBC leads the entire corpus with roughly 2,290 articles**, followed by the *Guardian*, Yahoo Finance, the *Daily Mail*, the NASDAQ newswire, Benzinga, Medscape, Investing.com, *Nature*, and the Motley Fool [news corpus]. A curated set of the leading business-and-markets outlets (CNBC, Yahoo Finance, the NASDAQ wire, Benzinga, Investing.com, the Motley Fool, MarketBeat, Bloomberg, Reuters-adjacent wires, and their peers) accounts for on the order of **6,300 articles, close to one in five in the whole corpus**, drawn from a long tail of 4,178 distinct publications [news corpus]. The peptide beat, in aggregate, is a business beat.

The topic tags confirm the framing. Across the corpus, "obesity" (22,853 articles) and "diabetes" (16,154) lead, as you would expect, but the very next tier is *commercial*: "regulatory" attaches to 14,633 articles, "market" to 14,399, "safety" to 13,302, and "supply chain" to 8,728 [news corpus]. Purely cultural and societal framings ("culture" 7,359, "society" 7,349) trail the market machinery. For a large share of readers, the peptide era did not begin with a doctor. It began with a chart of two share prices.

Those two prices belong to a *duopoly*. The GLP-1 franchise is, commercially, a two-company race: **Novo Nordisk** (semaglutide, liraglutide) and **Eli Lilly** (tirzepatide, and the investigational retatrutide and orforglipron), with a second tier of challengers such as Boehringer Ingelheim's survodutide circling the edges. The corpus tracks 659 companies in total, but coverage concentrates overwhelmingly on those two names, because the market structure genuinely is that concentrated [news corpus]. That concentration is itself the investment thesis and the investment risk in one: enormous, defensible franchises, and a scoreboard on which every trial readout for a next-generation molecule moves a mega-cap valuation. It is the rare therapeutic area where a Phase 2 body-weight number is front-page financial news.

NO REVENUE FIGURES HERE, ON PURPOSE

This report models media, commerce, and culture, not audited pharmaceutical sales. It deliberately reports no company revenue, market-size, or prescription-volume numbers, because it does not hold the administrative data required to state them responsibly. What the corpus supports is the *shape* of the business story, its financial-press dominance, its duopoly structure, its supply dynamics, and that shape is stated; the dollar magnitudes are left to sources built to measure them.

The event that turned the business story into a lived one for millions of patients was a *shortage*. From roughly 2023 into 2025, demand for GLP-1 drugs outran the ability to manufacture them, and the fallout is written all over the corpus. "Supply chain" is one of the largest topic tags in the entire dataset (8,728 articles), and two adjacent, revealing tags rise directly out of the shortage: "compounding" attaches to about **2,566 articles** and "telehealth" to about **2,214** [news corpus]. The mechanics matter. When a drug is officially in shortage, U.S. rules open a window for *compounding pharmacies*, which mix their own versions of the active ingredient to fill the gap, and a wave of *telehealth* clinics sprang up to prescribe and ship them. A drug-supply problem thus became a distribution revolution: it pulled GLP-1 prescribing out of the traditional clinic and into a fast, online, direct-to-consumer channel, and, critically, it normalized the idea of ordering an injectable peptide over the internet. That normalization is the bridge, examined in Part III, between the regulated GLP-1 world and the gray-market peptide world that grew up beside it. The shortage did not just strain supply. It taught a mass audience a new way to buy.

PART III

The other peptide boom

07 · Two media, two worlds

The central discovery of this study is not about any single molecule. It is about a *gap between two mirrors*. Point one mirror at the news and another at social media, ask each "what is a peptide," and they return almost opposite pictures. This is measurable, and the measurement is stark. When every article and every social post that can be tied to a specific compound is labelled by whether it names a GLP-1 drug or some other peptide, the two media split almost cleanly in half along the same seam, in opposite directions [news corpus], [social listening]:

WHERE THE CONVERSATION HAPPENS	COMPOUND-LINKED ITEMS	MENTION A GLP-1 DRUG	MENTION A NON-GLP-1 PEPTIDE
News corpus	32,659	~92.6%	~20.4%
Social corpus	8,242	~7.6%	~80.1%

Shares exceed 100% within a medium because one item can mention both a GLP-1 and a non-GLP-1 compound. Source: analysis.json glp1_control, computed over compound-linked items only (news: 30,226 GLP-1-linked of 32,659; social: 6,603 non-GLP-1-linked of 8,242). The contrast, not the decimal, is the point.

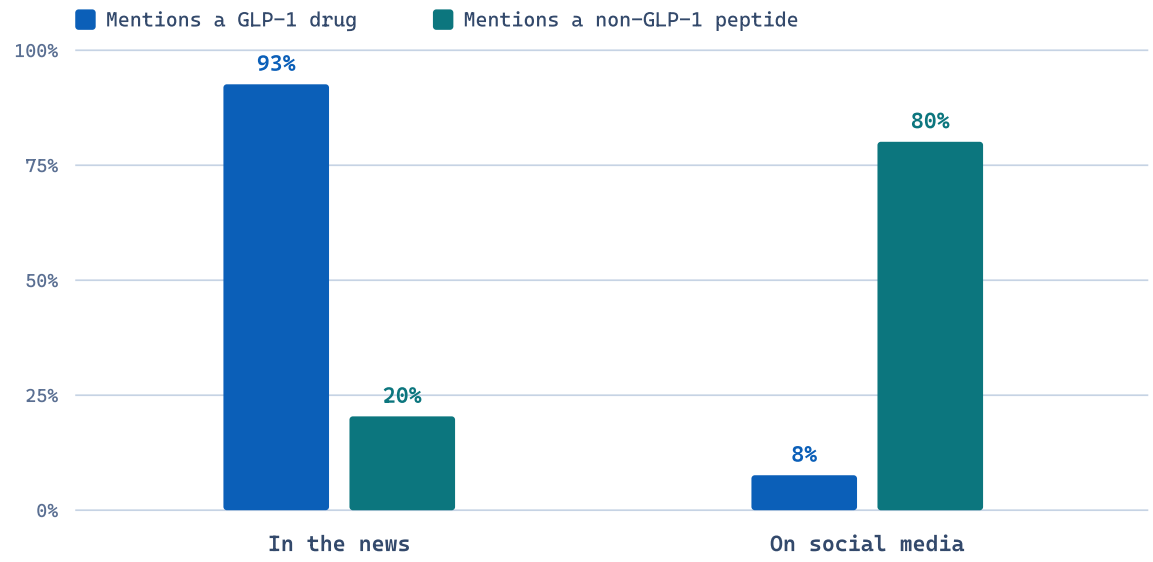

FIGURE 6 The two-media split. Share of compound-linked items that mention a GLP-1 drug versus a non-GLP-1 peptide, in the news corpus and the social corpus. The two media are near-mirror images. Percentages can exceed 100 within a medium because one item may mention both. Source: study corpus.

Sit with those numbers. In the news, **more than nine of every ten** compound-linked articles name a GLP-1 drug. On social media, **fewer than one in twelve** does, and four in five name something else entirely [social listening]. The press and the public are not covering the same subject at different volumes. They are covering different subjects.

This is why the peptide-boom question has been so badly mis-answered. Journalists have asked "is there a peptide boom beyond Ozempic?", looked at their own coverage, found it is roughly 93% GLP-1, and concluded, reasonably from where they sat, not really. But the coverage is not the culture. On social platforms, GLP-1 drugs are close to a rounding error, and the conversation is carried by a completely different cast: **GHK-Cu, BPC-157, TB-500, MOTS-c, GHK, KPV, IGF-1, ipamorelin, semax, epithalon** [social listening]. In the social share-of-voice ranking, the copper skincare peptide GHK-Cu (299 mentions) and the "healing" peptide BPC-157 (274) sit *above* semaglutide (286) and tirzepatide (244), the two best-selling drugs on earth [social listening]. The research peptides are not a fringe of the online conversation. They are its main body.

The reason the press missed it is structural, not lazy. Newsrooms are built to detect *events*: an FDA approval, an earnings surprise, a lawsuit, a celebrity confession. The GLP-1 story is a machine that produces those events on schedule, which is exactly why 4,178 publications cover it. The research-peptide world produces almost none of them. It has no public company to move, no approval to announce, no quarterly number to react to. What it produces instead is quieter and, to a newsroom's

instruments, nearly invisible: hundreds of millions of TikTok views, thousands of Reddit threads, a steadily lengthening roster of online vendors. There are, in effect, **two peptide booms running in parallel that barely touch**, and only one of them trips the sensors the mainstream press relies on. The rest of Part III is a map of the boom that does not.

08 · A field guide to the research peptides

If the GLP-1 drugs are one crowded row of the marketplace, the "research" peptides are the rest of the floor, and they are not a single thing. The underlying research catalog tracks **301 compounds** grouped by *function* rather than chemistry, a market taxonomy that mirrors how these molecules are actually sold, stacked, and discussed rather than how a chemist would file them [news corpus]. The largest functional groups map onto a handful of recognizable communities, each with its own goals, vocabulary, and archetypal user.

FUNCTIONAL FAMILY	COMPOUNDS TRACKED	REPRESENTATIVE MEMBERS	WHO REACHES FOR IT
Growth-hormone axis	20	Ipamorelin, CJC-1295, sermorelin, tesamorelin	Anti-aging and physique users chasing GH-like effects without GH
Nootropic / neuro	28	Semax, selank	Biohackers and students seeking focus and "brain fog" relief
Repair / recovery / muscle	16+	BPC-157, TB-500, IGF-1	Injured lifters and athletes seeking faster healing
Cosmetic / skin-hair	18+	GHK-Cu, GHK, "GLOW" blends	Skincare and hair-regrowth users, often crossing over from topical to injected
Bioregulator / longevity	20+	Epithalon, the Khavinson "Cytomax" family	Longevity and anti-aging adopters, often older
Metabolic / performance	25	Semaglutide, tirzepatide, retatrutide, 5-Amino-1MQ, MOTS-c	Weight-loss and body-composition seekers

Source: analysis.json taxonomy (catalog_group field of the 301-compound catalog). Figures are the largest single named groups; the catalog splits into many smaller hybrid categories (for example "Mitochondrial/longevity," "Repair/vascular"), which are consolidated here into the six functional families named in the text. This is a market taxonomy, not a pharmacological one; it deliberately includes non-peptides (SARMs, the nucleotide NAD+, the small molecule 5-Amino-1MQ) because they are sold and used inside the same subculture.

Walk the families in turn, because each answers a different human want. The *growth-hormone axis* peptides (ipamorelin, CJC-1295, sermorelin) are *secretagogues*: they coax the body's own pituitary into releasing more growth hormone, and they are marketed to people who want the recomposition and recovery benefits associated with GH without injecting GH itself. The *repair and recovery* group is the emotional heart of the movement: *BPC-157* (a fragment of a stomach-protective protein) and *TB-500* (a fragment of the protein thymosin beta-4) are sold on the promise of accelerated healing of tendons, joints, and injuries, and they draw the most fervent testimony of any category. The *cosmetic* peptides, led by *GHK-Cu* (a copper-bound tripeptide with a genuine, decades-old history in topical skincare) are the crossover story: users who first met the molecule in a face serum increasingly inject it in pursuit of skin, hair, and "glow" effects. The *nootropic* wing (*semax*, *selank*, both developed and approved in Russia, neither approved in the U.S.) belongs to the cognitive-enhancement and biohacking scene, a distinct subculture with its own logic. The *bioregulators* (*epithalon* and the Russian "Khavinson" peptide family) are the longevity fringe, short peptides claimed to act on aging itself, resting almost entirely on older non-English literature. And the *metabolic* family is where the two booms actually touch: it holds both the GLP-1 drugs and research compounds like *MOTS-c* (a mitochondrial-derived peptide) and *5-Amino-1MQ*, sold to the same weight-and-performance audience.

> THE DENOMINATOR PROBLEM, RESTATED
>
> Almost none of these families has the clinical evidence its market presence implies. BPC-157 has about 228 lifetime publications and two genuine human trials; GHK-Cu, one genuine trial; MOTS-c, one; semax and epithalon, zero registered trials at all [PubMed], [ClinicalTrials.gov]. What the taxonomy actually maps is not six bodies of medical evidence. It is six *communities of belief*, each organized around a want (youth, recovery, focus, beauty, longevity, leanness) and a molecule they have decided answers it.

09 · Two roads to prominence

The most revealing thing about a peptide is often not how much attention it has, but the *shape of how it got there*. Two products can reach the same level of prominence by opposite routes, and the route is diagnostic. Reconstruct each compound's timeline across the two media, when it first appears in the news versus in social posts, and the field sorts into two populations that took two roads.

The first road is *top-down*. A pharmaceutical company runs trials, results are published, the business press covers them, and public discussion follows the institutions. The second is *bottom-up*. A fitness or biohacking community starts experimenting, the discussion spreads on TikTok and Reddit, "research use only" vendors list the compound, and only much later, if ever, does a newsroom notice. The first road is pushed by regulators and pharma; the second is pulled by consumers running years ahead of both. Because social posts carry real timestamps back to 2013, while the news corpus can only be trusted for *ordering* rather than raw monthly volume (it was harvested in 2026, which inflates recent counts), the first-appearance comparison is the cleanest signal this dataset offers about which road a compound travelled [social listening], [news corpus].

COMPOUND	FIRST SOCIAL POST	FIRST NEWS ARTICLE	ROAD
Semaglutide	2019-06	2019-06	Top-down (simultaneous)
Tirzepatide	2022-11	2021-08	Top-down (news led)
Retatrutide	2024-04	2022-10	Top-down (news led, pre-approval)
Orforglipron	2024-09	2023-06	Top-down (news led)
BPC-157	2019-11	2024-06	Bottom-up (~4.5-year lag)
GHK-Cu	2020-07	2025-04	Bottom-up (~4.75-year lag)
MOTS-c	2020-07	2025-04	Bottom-up (~4.75-year lag)
TB-500	2021-03	2024-06	Bottom-up (~3-year lag)
Semax	2019-07	2025-06	Bottom-up (~6-year lag)
KPV	2020-07	none	Social-only

Source: analysis.json diffusion (first-appearance month in each medium, by compound). "Lag" is the gap between the first social post and the first news article. GLP-1 drugs appear in the news at or before their social debut; research peptides appear in social years before any news coverage. Ordering is reliable; raw monthly magnitudes are not, per the corpus harvest-recency caveat.

The pattern is unambiguous. Every GLP-1 drug reached the news at or before it reached social media. Every research peptide reached social media years before it reached the news, by three years for TB-500, closer to five for GHK-Cu and MOTS-c, and a remarkable six for semax. KPV never reached the news at all: 154 social mentions and zero articles, a compound that exists as a culture with no press coverage whatsoever [social listening]. When journalism arrives years after the conversation, it is not *amplifying* an official event. It is *reflecting* a phenomenon that formed without it, and that reversal is the single most reliable tell that a compound travelled the bottom-up road.

BPC-157 is the crossover case, the compound caught in the exact act of jumping from one world to the other. It appears in the social record from **November 2019**, built a large and devoted following through the early 2020s on the strength of injury-recovery testimony, and yet is essentially absent from the news corpus until **June 2024**, a lag of roughly four and a half years between culture and coverage [social listening], [news corpus]. What broke the silence was not a trial. It was the compound's own online popularity and the regulatory attention that popularity finally attracted, the U.S. FDA moved to restrict BPC-157 in the compounding context, which is precisely the kind of *event* a newsroom is built to cover. Now both curves are climbing together, with news volume cresting in 2026 as journalists and regulators engage at last. BPC-157 is what a research peptide looks like at the moment the mainstream catches up to it: still unapproved, still evidence-thin, but no longer

invisible. Its first controlled Phase 2 efficacy trial began only in **2026**, years after millions of doses had already been sold and self-administered on the strength of animal data and anecdote [ClinicalTrials.gov]. That inversion, use first, evidence maybe later, is the defining rhythm of the entire research-peptide world, and BPC-157 is its clearest single illustration.

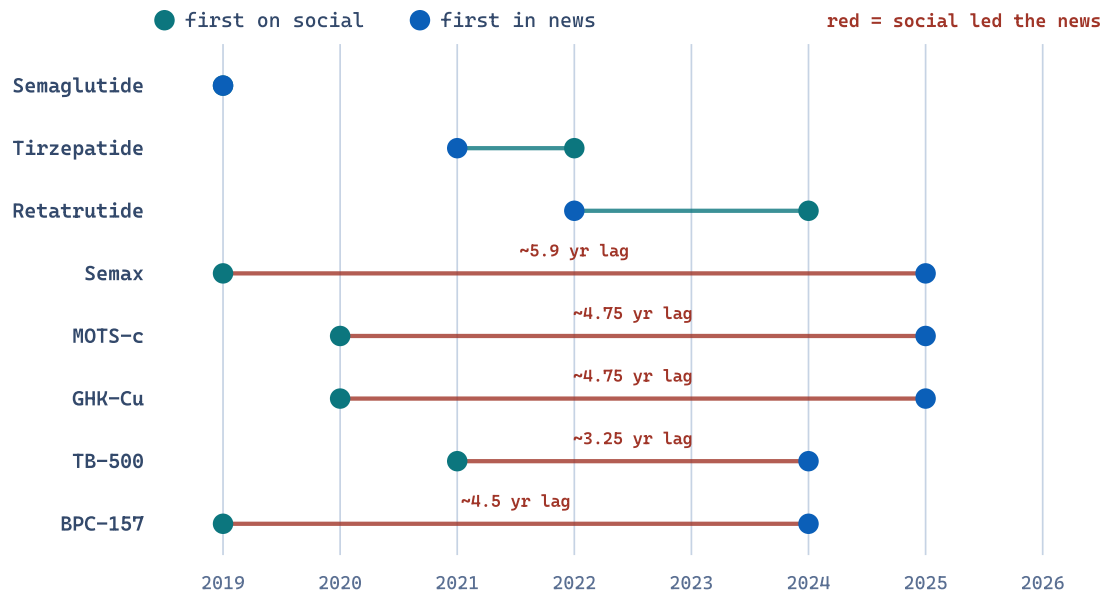

FIGURE 7 Two roads to prominence. The year each compound first appeared on social media (teal) versus in the news (blue). For the research peptides, social discussion led coverage by years (red); for the GLP-1 drugs, journalism arrived with or ahead of the culture. Source: study corpus.

The texture of the bottom-up road is visible in the discourse itself. The research-peptide conversation is dense with the specifics of actual use, dosing tables, reconstitution questions, week-by-week logs, and it reads nothing like press coverage. A BPC-157 user posts a structured before-and-after: "pain before: 4~6/10, pain now 2~3/10 (50%+ reduction)" [social listening]. A GHK-Cu user reports, of hair regrowth, "I can see some baby hairs growing back which is progress" [social listening]. A semax user describes "clarity, the mental static dropping, easier task initiation" [social listening]. These are anecdotes, evidence of *discourse*, never of efficacy, and this report keeps that line bright: "people are reporting recovery from BPC-157" and "BPC-157 has been shown to speed recovery in humans" are different sentences, and only the first is supported here. But as a record of how a molecule travels without a marketing department, of a culture teaching itself to use an unapproved drug in public, years before any institution looks up, the peptide internet has no equal. The press was late because it was watching the wrong road.

PART IV

The evidence, class by class

Part Three established the report's central structural finding: the news is a GLP-1 medium, the internet is a research-peptide medium, and the two booms barely overlap. Part Four takes each evidence class in turn and works it to the bottom, because the whole discipline of this report rests on never letting one class stand in for another. The order runs from the strongest kinds of evidence for establishing what is *true* of the world, the scientific literature and the clinical-trial record, down through what can and cannot be seen about actual use, and then into the commercial, search, social, journalistic, and advertising layers, where volume is abundant and proof is not. Each section closes with the read a professional should take from it.

10 · The scientific record

Before a compound is a market or a meme, it is, in principle, a question that science has or has not answered. Two public registries let us grade the answer directly. PubMed indexes the world's biomedical literature; queried compound by compound, it measures how much formal research a molecule has attracted. The picture it draws of the peptide field is not the picture the marketplace draws, and the divergence is the foundation for everything that follows.

Start with the scale of the gap. The nine GLP-1 and incretin compounds tracked here hold roughly **20,000 PubMed-indexed publications between them**, and the literature is not merely large, it is accelerating. In 2024 alone those nine compounds generated about **2,500 papers**, led by semaglutide with **1,052 in that single year**, liraglutide with **567**, and tirzepatide with **399** [PubMed]. Put the anchor years side by side and the trajectory is unmistakable: semaglutide went from **7 papers in 2015** to **160 in 2020** to over a thousand in 2024, and tirzepatide from **zero in 2015** to **11 in 2020** to nearly four hundred [PubMed]. This is the signature of a field into which a wave of successful trials and commercial money has poured resources. It is what a genuine, well-capitalized scientific boom looks like from the outside.

Now set the research peptides that dominate social media against that curve, and the contrast is severe.

COMPOUND	LIFETIME PUBMED RECORDS	2024 RECORDS	CHARACTER OF THE LITERATURE
Semaglutide	5,465	1,052	Exploding, human-trial-heavy
Liraglutide	5,868	567	Mature, clinical
Tirzepatide	2,391	399	Steep recent climb
Kisspeptin-10	950	34	Academic reproductive physiology

COMPOUND	LIFETIME PUBMED RECORDS	2024 RECORDS	CHARACTER OF THE LITERATURE
Thymosin alpha-1	864	14	Approved-drug (ex-US) oncology/hepatitis field
TB-500 / thymosin beta-4	1,073	15	Declining; mostly a distinct pharma program
MOTS-c	250	36	Small but the fastest-growing niche peptide
Melanotan II	243	7	Old dermatology/pharmacology base
Semax	231	4	Largely older non-English literature
BPC-157	228	7	Overwhelmingly preclinical/animal
GHK-Cu	189	9	Cosmetic-chemistry base
Epithalon	142	1	Sparse; older Russian work
Selank	135	3	Sparse; Russian-market anxiolytic
Ipamorelin	54	2	Tiny
CJC-1295	33	0	Negligible
AOD-9604	16	0	Negligible
5-Amino-1MQ	3	0	Effectively unstudied

Source: NCBI PubMed via direct query, as of August 2026 [PubMed]. Lifetime totals are publication-activity counts, not a measure of quality or human relevance; 2024 is the most recent full anchor year. NAD+ is excluded from this table because the term matches a ubiquitous cellular cofactor (about 85,700 records across all of biochemistry) and cannot be compared to a defined drug.

Three features of the research-peptide column deserve emphasis, because each undercuts a claim the marketplace makes implicitly. First, **the flagship consumer peptides have tiny lifetime literatures.** BPC-157, arguably the single most-hyped healing peptide online, has accumulated about **228 papers in its entire history and only around seven in 2024** [PubMed]. That is not the corpus of a validated therapy; it is the corpus of a compound that a handful of laboratories, several of them linked to one Croatian research group, have studied mostly in rats. Second, **some of these literatures are shrinking, not growing.** Thymosin beta-4, the molecule behind consumer "TB-500," has a respectable **1,073 lifetime papers, but the annual output fell from about 51 in 2015 to 15 in 2024** [PubMed], and most of it concerns a distinct ophthalmic and cardiac pharmaceutical program rather than the systemic recovery injection sold to consumers. A declining literature is the opposite of an emerging science. Third, **the largest non-GLP-1 literatures are not consumer-peptide**

literatures at all. Kisspeptin (about 950 papers) is an academic reproductive-endocrinology field, and thymosin alpha-1 (about 864) is an approved immunomodulator studied heavily in China for hepatitis and oncology [PubMed]. Neither reflects the wellness market that has adopted the injectable-peptide category.

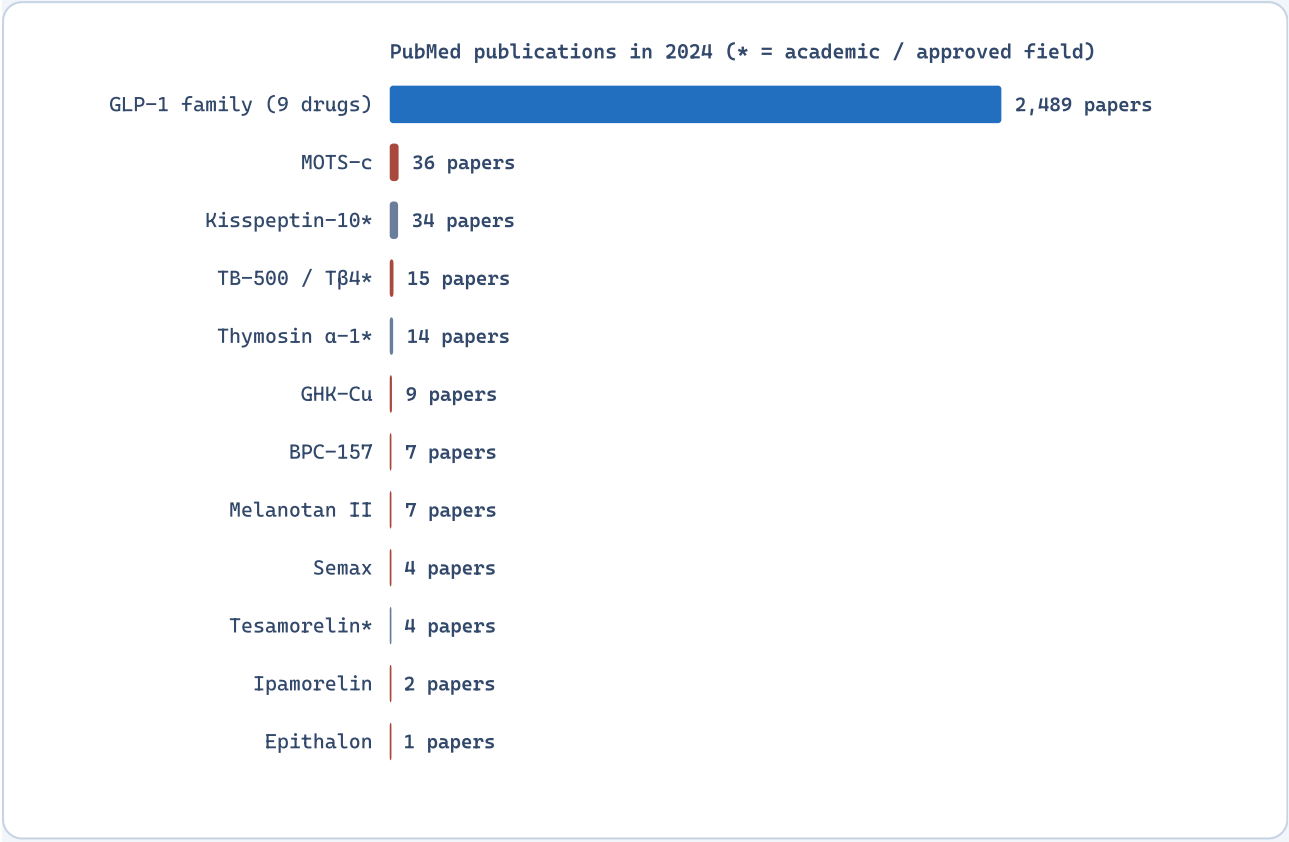

FIGURE 8 The literature gap. PubMed publications in 2024: the nine GLP-1 drugs together produced roughly 2,500 papers, while the most-discussed research peptides produced a handful each. Starred entries are academic or approved-drug fields, not consumer-peptide interest. Source: PubMed.

READ THE COUNTS CAREFULLY

Publication volume measures **activity, not truth**. A compound with 200 papers is not "proven": much of any research peptide's literature is animal work, cell-culture studies, and review articles that recycle the same preclinical findings rather than new human data. Conversely, a low count does not always mean no knowledge, some Russian-developed peptides (semax, epithalon, selank) have older non-English literatures that PubMed indexes poorly, which is why this report reads the numbers comparatively and directionally rather than as a scoreboard. The honest summary is that the scientific establishment is investing enormously in GLP-1 biology and comparatively almost nothing in the compounds the wellness market has embraced.

The investor and professional read. Publication activity is the earliest and cheapest leading indicator of where durable, defensible value is being created, and it points almost entirely at the incretin class. For a compound whose entire commercial case rests on biological effect, a flat or shrinking literature is a red flag that the science is not catching up to the marketing, not a neutral fact. The practical asymmetry to carry forward: on one side of this market sit molecules with thousands of studies and a vertical research curve; on the other sit products with a few hundred mostly-preclinical papers and, in several cases, single digits. Any thesis that treats "peptides" as one investable category is blurring a distinction the literature draws in the sharpest possible terms.

11 · Clinical development

Publications tell you a compound is being studied. Clinical-trial registrations tell you it is being tested in *people*, which is the threshold that matters for anyone weighing whether to use it, sell it, or underwrite it. ClinicalTrials.gov, the NIH registry of human studies, lets us count that threshold compound by compound, and here the divide is starker than in the literature. It also produces what this report considers the single most important fact about the research-peptide market.

Human testing is concentrated almost entirely in the GLP-1 class. The registry lists roughly **741 trials for semaglutide, 513 for liraglutide, 377 for exenatide, and 274 for tirzepatide**, with dozens to hundreds more across dulaglutide (139), thymosin alpha-1 (63, discussed below), orforglipron (50), cagrilintide (43), retatrutide (33), and survodutide (26). The nine incretin compounds together carry about **2,196 registered trials** [ClinicalTrials.gov]. Semaglutide and liraglutide *alone*, at 1,254 trials, outnumber every non-GLP-1 research peptide in this study combined by roughly an order of magnitude. These are mature, multi-phase programs, the kind that underpin approved medicines and that cost hundreds of millions of dollars to run.

Among non-GLP-1 peptides, the only well-developed ones are the already-approved ones. Thymosin alpha-1 (thymalfasin, approved as Zadaxin in many countries) carries about **63 trials**, weighted toward Phase 2 and heavily Chinese, in hepatitis, sepsis, and oncology-adjuvant settings. Tesamorelin (approved as Egrifta for HIV-associated lipodystrophy) carries about **24 trials**. Bremelanotide, sold as PT-141 online and approved as Vyleesi, carries about **10 trials** spanning a full Phase 1 to Phase 3 program, which makes it the **most clinically validated of the entire non-GLP-1 "research" set** [ClinicalTrials.gov]. Approval and clinical evidence travel together, as they should. The lesson is that the non-GLP-1 peptides with real trial records are precisely the ones that went through a regulator somewhere, not the ones the internet discovered.

The compounds the internet loves have almost no human evidence at all. Counting only genuine, compound-specific interventional trials, the picture is close to empty:

COMPOUND	REGISTERED TRIALS	GENUINE COMPOUND-SPECIFIC	STATUS OF THE EVIDENCE
BPC-157	3	2	One 2015 Phase 1, one 2026 Phase 2 (recruiting)
GHK-Cu	3	1	One topical Phase 2 beginning 2026

COMPOUND	REGISTERED TRIALS	GENUINE COMPOUND-SPECIFIC	STATUS OF THE EVIDENCE
MOTS-c	5	1	One Phase 2a beginning 2026; rest observational
Ipamorelin	3	2	Two Phase 2s from a program abandoned c. 2013
CJC-1295	1	1	A single 2005 Phase 2, terminated
Melanotan II	1	1	One 2026 vitiligo Phase 2
Semax	0	0	No registered trials
Epithalon	0	0	No registered trials
Selank	2	0	Both matches unrelated (tDCS/rTMS studies)
AOD-9604	0	0	No trials under this term
5-Amino-1MQ	0	0	Preclinical only; not a peptide

Source: ClinicalTrials.gov API v2, intervention-name search, as of August 2026 [ClinicalTrials.gov]. "Registered trials" is the raw count the registry returns; "genuine compound-specific" strips out synonym-expansion false positives and name collisions (for example, the raw "selank" search returns unrelated brain-stimulation studies, and "sermorelin" expands to the whole growth-hormone-releasing-hormone class). Zero registered trials does not prove zero human exposure; for several Russian-market compounds it means the exposure happened outside the trial system entirely.

Add it up and the entire trendy injectable-peptide category, the compounds filling the social feeds and the vendor catalogs, totals **well under thirty genuine registered trials**, and almost none are completed efficacy studies. Contrast that with a single incretin drug's 741. The market that has formed around these molecules did not form on the strength of human data, because there is almost none.

THE CLINCHER, ONE SPONSOR

Look closely at the recent trials and a remarkable pattern appears. Nearly every *first real* interventional study of BPC-157, GHK-Cu, MOTS-c, melanotan II, and the genuine TB-500 (thymosin beta-4 17-23) fragment is a single Phase 2 registered by **one company, Hudson Biotech, all beginning in February 2026** [ClinicalTrials.gov]. Before that lone sponsor stepped in, most of these compounds had *no* company-run efficacy trial in human history. Millions of doses

> had already been sold and self-administered on the strength of animal data and anecdote; the first controlled human tests of whether they work are only starting now, years *after* the market and the culture formed around them.

That inversion is the defining structural feature of the research-peptide phenomenon, and it is worth naming precisely: **use first, evidence maybe later.** In the textbook pharmaceutical sequence, controlled human evidence is the gate that adoption must pass through; a drug is studied, then approved, then prescribed, then used. In the research-peptide world the sequence runs backwards. Adoption came first, built on preclinical signals and self-reported experience, and the evidence, if it arrives at all, arrives to a market that has already decided. The Hudson Biotech trials are the first attempt to run the sequence forward. Their results, due over the next few years, will either begin to validate a handful of these compounds or begin to falsify them, and either way they represent the first time the question has been asked under controlled conditions.

The investor and professional read. The trial record is the cleanest available proxy for regulatory risk and for the probability that a compound's efficacy claims survive contact with controlled data. On that measure the research peptides are not early-stage assets with thin pipelines; most are pre-clinical assets with a consumer market attached, which is a fundamentally different and less stable configuration. The concentration of first-in-human efficacy work in a single 2026 sponsor is a genuine inflection point to watch: it is simultaneously the first evidence catalyst the category has ever had and a single point of failure. A clean positive readout could pull one or two compounds toward legitimization and reprice the category; a null or safety-flagged readout could hand regulators exactly the evidence a crackdown would need. Either outcome matters more to the research-peptide market than any amount of the social volume the next sections measure.

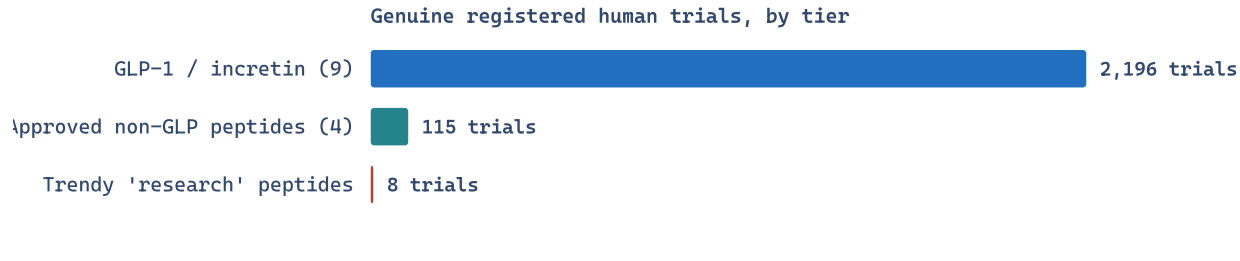

FIGURE 9 Use first, evidence later. Genuine registered human trials by tier: the incretin class dwarfs the approved non-GLP peptides, which in turn dwarf the trendy research peptides that fill the vendor catalogs. Source: ClinicalTrials.gov.

12 · What we can and cannot see about real-world use

Between "a compound exists" and "people are talking about it" sits the question that matters most for public health and is hardest to answer here: how many people are actually using these things, and how? For approved medicines that question has authoritative answers, prescription claims, pharmacy dispensing data, insurer records, national health surveys, and for the GLP-1 drugs those sources exist and describe tens of millions of users. This study does not have direct access to that administrative layer, which sits at the very top of the evidence hierarchy, and the honest move is to say so rather than substitute a noisier signal and dress it up as a count.

For the research-peptide world the answer is starker still: **no reliable population estimate of use exists anywhere.** These compounds are sold as laboratory chemicals, bought without prescriptions, and used without medical oversight, which means there is no claims database, no prescriber writing a script, and no survey instrument built to count them. The structure that makes them legally available is exactly the structure that makes them statistically invisible. Anyone who asserts that "X million Americans are using BPC-157" is guessing, because the data that could support such a number does not exist. This is not a gap this report can close, and it does not pretend to.

What *can* be observed are the indirect footprints of use, the visible edges of an activity that is mostly hidden.

- **A commercial supply that exists only because demand exists.** The price tracker alone catalogs thousands of live product listings across dozens of active U.S. sellers (next section). Vendors do not carry, restock, and price-compete on inventory that nobody buys; a competitive supply is itself evidence of a real, recurring customer base, even though it cannot be converted into a headcount.
- **A discourse dense with the texture of genuine use.** The social corpus is full of dosing protocols, reconstitution questions, "week three update" posts, injection-site reports, and specific milligram figures. People describing a compound's sting on injection, a vial's cloudiness after mixing, or a precise weekly schedule are very probably using it, even though a single post proves nothing on its own (the social section returns to this).
- **Adjacent regulated channels the news corpus does track.** Telehealth clinics and compounding pharmacies are the visible, regulated cousins of the gray market, and they leave a media footprint. "Compounding" attaches to roughly **2,600 articles** and "telehealth" to about **2,200** in the news corpus [news corpus], a steady drumbeat reflecting how the GLP-1 shortage of 2023 to 2025 pulled peptide prescribing into a fast-growing online-clinic channel. That channel's growth is the closest thing to a hard signal that *supervised* peptide use has expanded.

The gap between these footprints and a real usage number is unbridgeable with the evidence in hand, and this report does not try to bridge it. It treats commercial availability and online discussion as strong evidence of a *market* and a *culture*, which they genuinely are, and refuses to convert them into a count of users, which they cannot support.

The investor and professional read. The measurement vacuum is itself a market fact with two edges. On one edge, the absence of utilization data means the true size of the research-peptide market is unknown and quite possibly larger than any visible proxy suggests, which is part of what makes it attractive to new entrants. On the other, that same opacity is precisely what regulators, payers, and litigators find intolerable, and it is the vulnerability a single adverse event could convert into decisive action. The regulated telehealth and compounding channel is the part of this world where usage is at least partially observable, which is one reason it is the most investable adjacency: it captures a share of the same demand while leaving an auditable trail. The gray market's invisibility is a feature for its participants and a latent liability for anyone exposed to it.

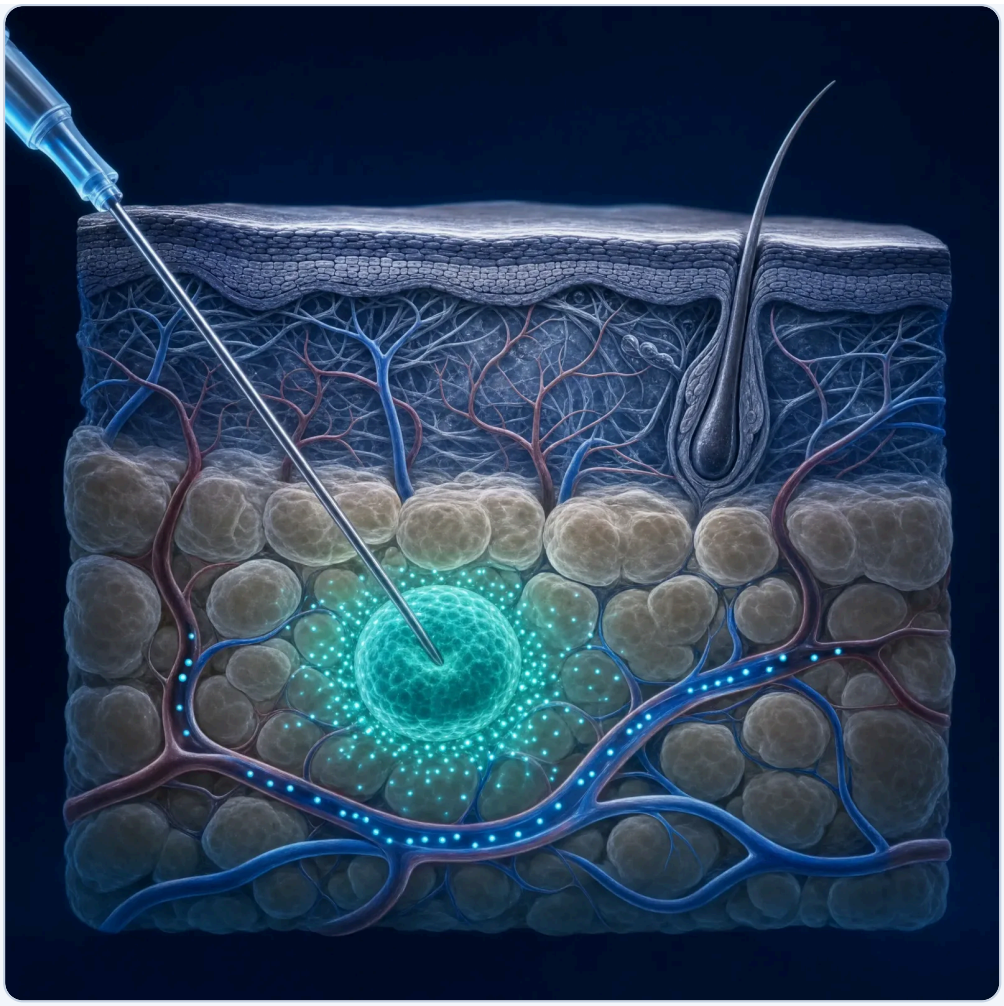

FIGURE 10 How these drugs enter the body: the subcutaneous depot. Most peptide therapies are injected into the fat just under the skin, forming a small depot that releases the drug slowly into the bloodstream over hours to days. Illustrative cross-section, not anatomically exact; it depicts a route of administration and is not a recommendation of any practice.

13 · The new marketplace

If the news is where the GLP-1 story lives, the marketplace is where the research-peptide story becomes concrete and countable. Running alongside this study is a live price tracker that crawls U.S. research-use-only peptide retailers under a strict, politely identified, no-evasion policy: it obeys robots rules and terms of service, and it excludes any storefront that is bot-walled or login-gated. What it captures is therefore a floor on the market, not a census, and even that floor is substantial. As of the August 2026 snapshot the tracker had catalogued more than **1,000 candidate seller domains**, confirmed **443** as verified retailers, and captured live pricing from **302** of them, spanning roughly **18,400 product listings** across **114 distinct compounds**, on about **27,800 individual price observations** [RUO price tracker]. The tracker is still expanding as it runs, so these figures should be read as a conservative snapshot of a growing market rather than its outer limit. Seller, listing, and compound counts here are current as of this pass; the price-per-milligram medians and dispersion statistics elsewhere in this section were computed on the July–August price-collection window.

The first thing the marketplace reveals is that **the research-peptide economy is real, broad, and priced like a competitive commodity market.** The clearest evidence is price dispersion. The same compound at the same vial size routinely varies several-fold from one seller to the next: a 5 mg vial of BPC-157, for instance, ranges from roughly **$22 to $250** across sellers in the corpus, and comparable spreads appear for MOTS-c, retatrutide, and the GLP-1 drugs [RUO price tracker]. That kind of dispersion is the signature of a young, fragmented market with many entrants, weak brand consolidation, and buyers who cannot easily verify what they are getting, so price competes without quality being pinned down. Dispersion this wide is not evidence of a rational market pricing quality; it is evidence of a market still discovering its own floor.

The second thing the marketplace reveals is how money and attention are distributed across categories. Grouping the priced listings by function and taking the median price per milligram:

CATEGORY	LIVE LISTINGS	MEDIAN $/MG
Growth & anti-aging (GH-axis)	740	$7.50
GLP-1 & weight management	435	$7.27
Repair & aesthetic (BPC-157, TB-500, GHK)	488	$6.40
Cognitive, aesthetic & ancillary	355	$5.00
Bioregulators (Khavinson family)	179	$3.25
Ancillary compounds	45	$1.33
Topical cosmetic peptides	45	$1.02

Source: RUO price tracker, median of the most recent in-stock price per milligram within each category, collection window July to August 2026 [RUO price tracker]. Listing prices before shipping; not a claim about purity, identity, or legality. Per-milligram figures are comparable only within roughly similar dose ranges, not across wildly different vial masses.

The GLP-1 drugs, which own the news, are just **one row here**, and neither the most stocked nor the most expensive one. The growth-hormone and repair/aesthetic peptides carry more listings and comparable or higher per-milligram prices. In the marketplace, in other words, the research peptides are not a fringe; they are the main floor, and the GLP-1 drugs are a single well-lit aisle within a much larger store.

That impression is reinforced by an independent market-strength model built for the companion project, which blends commercial shelf presence and search-intent signals into a 0 to 100 score [P08 market model]. Its top ranks are striking for who is and is not there:

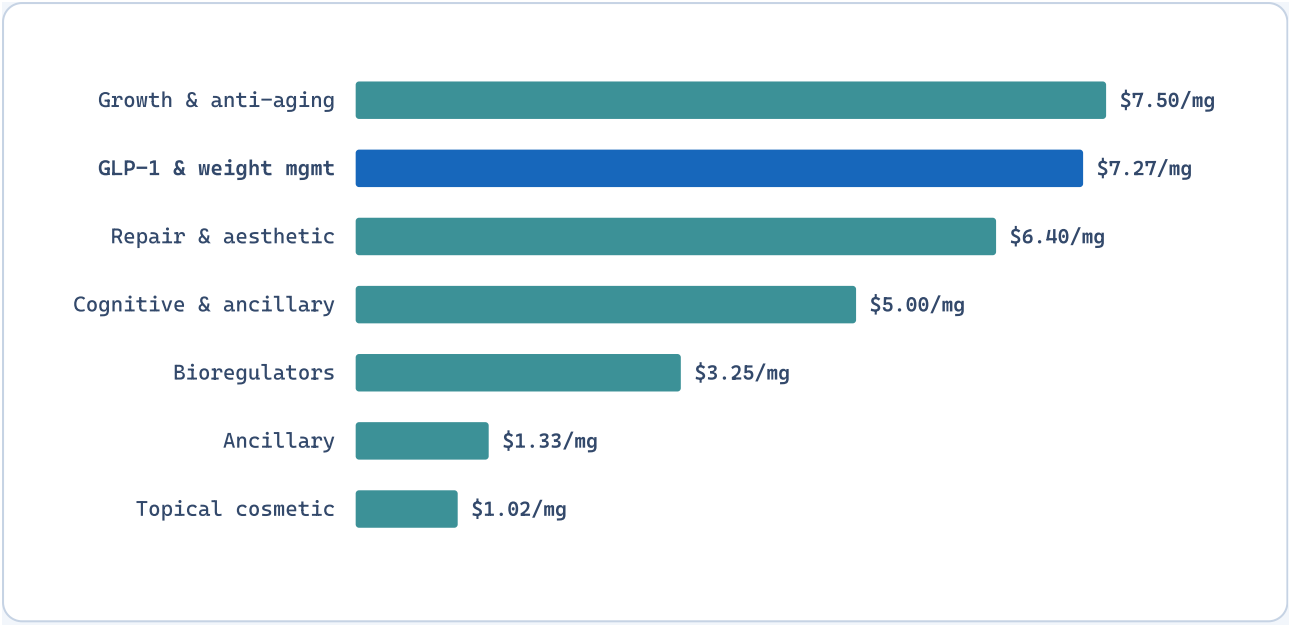

FIGURE 11 What each category costs. Median advertised price per milligram by functional category across live U.S. research-use-only listings, July to August 2026. The GLP-1 class is one row among several, not the most expensive. Listing prices before shipping; not a claim about purity or identity. Source: RUO price tracker.

RANK	COMPOUND	MARKET-STRENGTH SCORE	CLINICAL STATUS
1	Retatrutide	98	Investigational (in trials)
2	Tirzepatide	94	Approved
3	Tesamorelin	94	Approved (niche indication)
4	MOTS-c	93	Research chemical
5	BPC-157	93	Research chemical
6	GHK-Cu	93	Cosmetic ingredient / research
8	Semaglutide	90	Approved
9	Epithalon	89	Research chemical

Source: P08 market-strength model, 0 to 100, blending commercial shelf presence and search-intent signals [P08 market model]. Rank 7 is a BPC-157/TB-500 blend, omitted here as a formulation rather than a distinct compound. "Clinical status" is regulatory reality, not part of the score.

THE TELL

> The commercial market ranks MOTS-c, BPC-157, and GHK-Cu *above* semaglutide, the best-selling drug on earth, and rates unapproved research chemicals as equal in "market strength" to approved medicines. The market is not pricing clinical evidence. It is pricing attention and availability. This is the single clearest sign in the entire study that commercial adoption has decoupled from scientific validation, and it is exactly the decoupling Sections one and two of this part quantified from the other direction.

A final layer the marketplace exposes is how thin its trust signals are. The tracker attempted to attach community-sentiment signals to sellers and mostly could not: of the **30 retailers checked** against public discussion, **18 returned no meaningful signal**, a handful showed mixed or negative sentiment, and **two carried an explicit caution for suspected review manipulation** (astroturfing) [RUO price tracker]. Claims of domestic manufacture are similarly sparse: of the priced sellers whose homepages were checked, about **60 said nothing at all about where their product is made**, and only around a dozen made any explicit USA-manufacture claim, most of it marketing copy rather than verifiable provenance [RUO price tracker]. A young market with weak reputational signals and opaque sourcing is precisely the environment in which purity and identity cannot be assumed, which is why this report treats "a product is sold" and "a product is what it says it is" as separate claims.

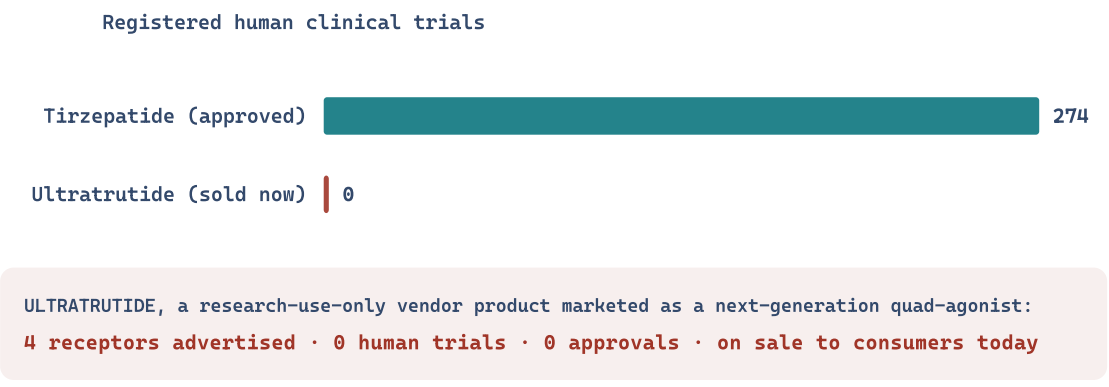

FIGURE 12 Attention without evidence. Ultratrutide is a real research-use-only vendor product marketed as a next-generation quad-agonist incretin. It has no registered human trials and no approval, yet is on sale today, the cleanest illustration of a market that prices attention rather than evidence. Trial counts: ClinicalTrials.gov.

The investor and professional read. The marketplace is where this phenomenon is most legible as a business, and its structure carries clear signals. Wide intra-compound price dispersion, dozens of undifferentiated sellers, and near-absent trust infrastructure describe a market in the land-grab phase, before consolidation, branding, and quality certification create durable margins. The opportunity that structure implies is a trust and verification layer, third-party purity testing,

verified provenance, reputation systems, because that is the scarce good the current market conspicuously lacks. The risk that same structure implies is that the whole edifice is legally contingent: it prices attention, not evidence or approval, so a regulatory reclassification or a high-profile purity scandal could reprice it abruptly. Anyone modeling this market should treat the market-strength ranking as a map of where demand currently points, not as a measure of which compounds have anything durable underneath them.

14 · Search and consumer curiosity

When people first hear about something, they search for it, which makes search behavior one of the earliest and most honest signals of genuine curiosity, the moment private interest becomes a measurable act, before anyone has bought or posted anything. This study carries a specific, disclosed limitation here: **it does not have direct access to Google Trends' official search-volume series.** Rather than dress up a substitute as the real thing, the report states the gap plainly and leans on two disciplined proxies instead.

The first proxy is the **search-intent component of the market-strength model** [P08 market model]. It is not a clean measure of public curiosity, it is weighted toward commercial keyword volume and therefore toward buying intent, but it consistently ranks the same metabolic and repair peptides at the top that the marketplace does. The convergence is itself informative: commercial search demand tracks the shelf, not the newspaper, which tells you that the people searching are, by and large, people looking to buy rather than people reacting to a news cycle.

The second proxy, and the more revealing one, is **social-media volume treated as a revealed-interest signal.** Curiosity exists on a ladder of effort. Typing a word into a search box is the cheapest rung. Posting about a compound, asking a dosing question, defending a protocol, or watching a ten-minute explainer to the end all sit higher on that ladder, because they cost more attention and often more social risk. On that measure the research peptides again dominate, and, as the next section shows, they dominate on the platforms where the effort per interaction is highest. Revealed interest of this kind is a stronger signal than a raw search count precisely because it is harder to fake and harder to do idly.

> WHAT THE DATA CANNOT TELL US
>
> Without the official search series, this report cannot say whether interest in a given peptide is rising or falling month to month, cannot cleanly separate "curiosity" from "intent to buy," and cannot benchmark peptide search volume against unrelated topics. Those questions stay open. Acquiring the trends series is the single highest-value addition a future edition could make; until then the report uses the signals it can actually verify and declines to infer a curve it cannot see.

The investor and professional read. The honest position is that consumer curiosity about peptides is broad and real, the marketplace and the social record both require it as an explanation, but its precise shape over time is the thing this evidence base is least equipped to measure. For anyone relying on this analysis, that is a reason to weight the verifiable commercial and social signals

heavily and to treat any confident month-to-month "search is surging" narrative with suspicion until an official series is in hand. The gap is also a concrete, cheap information edge available to whoever acquires clean search data first: it would convert the report's directional claims about curiosity into a dated, comparable time series, which is exactly what is missing.

15 · The peptide internet

If the non-GLP-1 peptide movement has a home, it is social media, and the collected social corpus, roughly **18,100 posts** tied to specific compounds, lets us look at that home directly. The first thing to understand is that "social media" is not one place. The three platforms in this corpus play three distinct roles, and the raw numbers make the division of labor visible at a glance.

PLATFORM	POSTS	TOTAL ENGAGEMENT	DISTINCT AUTHORS	ROLE
TikTok	11,688	~312 million	9,748	Reach and performance
Instagram	4,583	~273 thousand	2,894	Aesthetics and before/after
Reddit	1,876	~1 thousand	1,379	Deliberation and detail

Source: collected social corpus, compound-linked posts [social listening]. Engagement is summed platform-native interactions (likes, comments, shares) as captured; the three platforms are not directly comparable on engagement because each counts and exposes it differently. Reddit's tiny engagement number reflects how the platform surfaces interactions, not how much conversation occurs there.

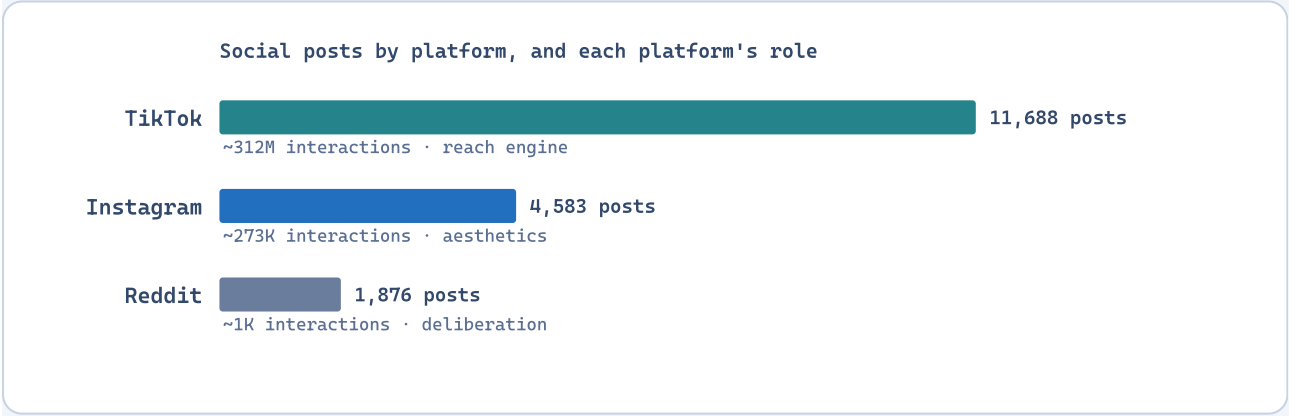

FIGURE 13 Three platforms, three jobs. Compound-linked social posts by platform. TikTok is the reach engine, Instagram the aesthetic showcase, Reddit the deliberation room. Engagement scales differ by orders of magnitude and are noted rather than plotted. Source: social listening.

TikTok is where peptides are performed. With about **312 million captured interactions across fewer than 12,000 posts**, it is overwhelmingly the reach engine, the surface where a compound goes from unknown to "everyone is talking about it," through transformation videos, day-in-the-life protocols, and confident explainers. Its cultural function is amplification, and its content is optimized for it. **Reddit is where peptides are deliberated.** Its engagement totals are tiny by comparison, but that undersells it badly: Reddit is where the detailed, technical, skeptical

conversation happens, the dosing tables, the sourcing debates, the "did anyone else get this side effect" threads. A Reddit post reaches few people but carries far more information per post, which is why it is the richest source for genuine experience data. **Instagram sits between them**, tilted toward the aesthetic and cosmetic end, skin, hair, anti-aging, consistent with the strong showing of GHK-Cu and other cosmetic peptides in its mix.

The cast of characters on social is the near-inverse of the news. Ranked by mention count, the most-discussed compounds are **GHK-Cu (299), BPC-157 (274), NAD+ (267), GHK (261), TB-500 (206), IGF-1 (189), MOTS-c (179), and ipamorelin (164)**, with the GLP-1 drugs semaglutide (286) and tirzepatide (244) present but not dominant [social listening]. These are healing, recovery, cognitive, cosmetic, and anti-aging compounds, almost none of them approved drugs. The vocabulary is its own dialect, and it is worth hearing directly, because the texture is the evidence. Compounds are *stacked*; protocols are *run*; vials are *reconstituted*; and the community polices *sourcing* with the seriousness of people who know they are buying unregulated products. A recurring tell of the culture is that users, keeping up the "research use only" fiction, narrate their own doses in the third person as happening to "my rat", one poster writes, "My rat is peaking physically and mentally," describing himself [social listening].

A handful of representative, clean posts convey what the discourse actually sounds like across the effect categories the movement organizes itself around:

- Recovery and repair, quantified in the community's characteristic self-tracking style: *"pain before: 4-6/10, pain now 2-3/10 (50%+ reduction); blood pooling before 6/10, now 2-3/10"* (BPC-157) [social listening].
- Performance, the fitness register: *"My run went through the roof. Smashed 6 PBs, Garmin's personalized VO2Max pace for the 4 intervals felt like a walk in the park"* (MOTS-c) [social listening].
- Cognitive optimization, the nootropic wing: *"What people report is clarity, the mental static dropping, easier task initiation, less procrastination"* (semax) [social listening].
- The connoisseur's vocabulary of formulations: *"No DAC feels more natural, pairs really well with Ipamorelin, and gives you way more control"* (CJC-1295) [social listening].
- And, importantly, honest ambivalence, which the hype narrative tends to omit: *"although I'm happy with the effects on appetite and weight, I have noticeably less drive and energy"* (retatrutide) [social listening].

These are not uniformly glowing. The corpus contains disappointment, side-effect reports (fever and joint pain after MOTS-c, unexpected hormonal shifts, injection-site swelling), and outright skepticism alongside the enthusiasm. That mix is part of what marks this as an established subculture with norms, jargon, and internal expertise rather than a wave of novelty. It is a community that has moved well past first contact into something with its own accumulated craft.

> DISCIPLINE CHECK
>
> Everything in this section is evidence of **discourse**, not of fact. That a compound is discussed intensely shows that people find it interesting, are experimenting with it, and are teaching each other how to use it. It does not show that the compound works, that it is safe, or how many people use it. "People are reporting remarkable recovery from BPC-157" and "BPC-157 has been shown to speed

> recovery in humans" are completely different sentences, and only the first is supported here. Every quote above is a data point about belief and behavior, never about biology.

The investor and professional read. The platform split is a distribution map. TikTok is the top of the funnel, the demand-creation engine that turns an obscure molecule into a searched, bought product; Reddit is the retention and expertise layer where committed users refine protocols and vet sellers; Instagram is the aesthetic on-ramp. A brand or clinic trying to reach this market has to understand that these are three different audiences at three different funnel stages, not one. Two structural facts should temper any read of the social volume, though. First, this is discourse, so it measures mindshare, not efficacy or safety, and mindshare in an unregulated category is unusually manipulable. Second, the same influencer-driven, sourcing-obsessed culture that builds demand also blurs the line between peer advice and paid promotion, which the advertising section returns to. Social attention is the leading indicator of where this market goes next; it is not, and must not be read as, evidence that any of it works.

16 · How journalism changed the conversation

Journalism did not create the peptide moment, but it shaped which parts of it the public sees, and the collected news corpus lets us characterize that coverage precisely rather than guess at it. Three features stand out, and together they explain why the mainstream record has been such an unreliable guide to what is actually happening.

The coverage is financial before it is medical. Sorted by outlet, the publishers covering peptides most heavily are not health desks but business-news operations. CNBC leads the entire corpus by a wide margin (about **2,290 articles**), followed by a cluster of financial and general outlets: Yahoo Finance (741), the NASDAQ newswire (464), Benzinga (384), Investing.com (377), the Motley Fool (313), MarketBeat (311), and Reuters (225), interleaved with consumer and lifestyle titles that cover the drugs as celebrity and body-image stories (the *Daily Mail* at 475, the *Guardian* at 767) and a thinner tier of genuine medical and scientific sources (Medscape at 378, *Nature* at 334) [news corpus].

OUTLET	ARTICLES	PRIMARY LENS
CNBC	2,290	Financial / markets
The Guardian	767	Consumer / lifestyle
Yahoo! Finance	741	Financial / markets
Daily Mail	475	Consumer / lifestyle
NASDAQ	464	Financial newswire
Benzinga	384	Financial / markets
Medscape	378	Medical / clinical

OUTLET	ARTICLES	PRIMARY LENS
Investing.com	377	Financial / markets
Nature	334	Scientific
The Motley Fool	313	Financial / investing

Source: news corpus, article counts by publisher [news corpus]. "Primary lens" is this report's characterization of each outlet's dominant frame, not a field in the data. The composition is the point: business and consumer-lifestyle outlets vastly outnumber medical and scientific ones.

For most of the public, in other words, the peptide story arrived framed as an *investment* story, Novo Nordisk and Eli Lilly share prices, shortage economics, patent fights, with the human and medical dimensions layered on top. That framing is not wrong, but it is partial, and it helps explain why the research-peptide subculture, which has no public company to cover and generates no earnings surprises, stayed invisible to the mainstream for so long. A newsroom built to cover markets sees the part of the peptide world that has a ticker.

A large share of the "coverage" is the same story repeated. Raw article counts overstate genuine attention because wire services and syndication multiply a single piece across many outlets. A duplicate-detection pass on the corpus groups roughly **6,200 articles into about 2,000 clusters**, split between exact duplicates (around 3,000 articles in 633 clusters) and looser syndication (around 3,100 in 1,365 clusters) [news corpus]. That is **nearly one in five articles in the entire corpus** contributing no independent reporting. The methodological lesson is general and it disciplines every count in this report: **syndication is not the same as independent attention**, and any measure of "how much was written" has to be discounted for the echo, which is why this report leans on de-duplicated and share-of-voice measures throughout rather than raw totals.

For different compounds, journalism led or lagged the public, and which one is diagnostic. For the GLP-1 drugs, the press *led*: coverage of semaglutide and tirzepatide began with their trials and approvals and ran ahead of, or alongside, public interest, the classic top-down path. For the research peptides, the press *lagged*, often by years. The cleanest case is BPC-157, which appears in the social record from **late 2019** but is essentially absent from the news corpus until **mid-2024**, a lag of roughly four to five years between the culture and the coverage [news corpus] [social listening]. GHK-Cu shows the same shape (social from 2020, news not until 2025), and some compounds, KPV is the example, have a substantial social presence and **no meaningful news footprint at all** [social listening]. When journalism arrives years after the culture, it is *reflecting* a phenomenon, not *amplifying* an official one, and that reversal is one of the most reliable tells that a compound traveled the bottom-up path rather than the top-down one.

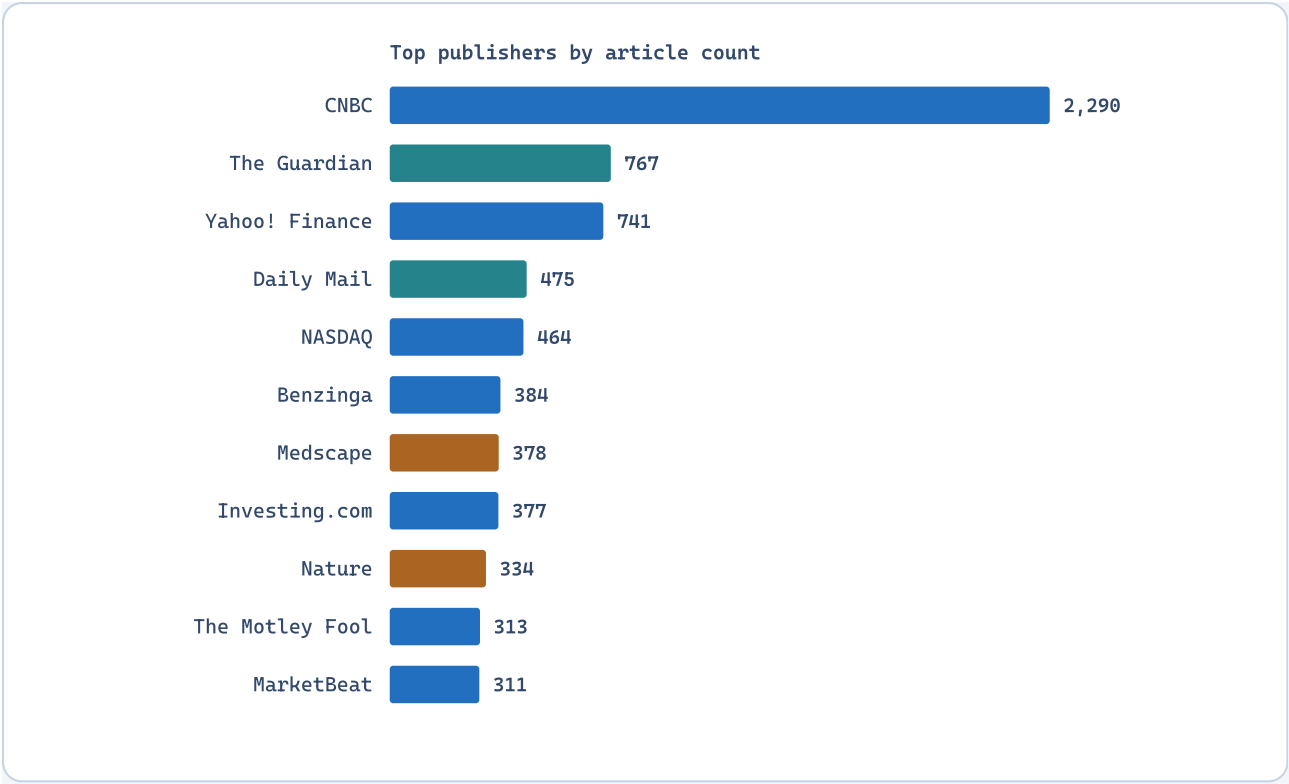

FIGURE 14 The press is a business desk. The outlets covering peptides most heavily are financial-news operations (blue), with consumer titles (teal) and a thinner medical tier (amber). The peptide story reached the public as an investment story first. Source: news corpus.

The investor and professional read. The composition of coverage is a map of where the information asymmetries sit. Because the press is structurally a business desk, it covers the peptide world with a ticker attached and is systematically late to, or blind to, the part that does not. For anyone trying to see an emerging consumer-health behavior early, mainstream coverage is a lagging indicator for exactly the compounds where being early matters most, and its raw volume is further inflated by syndication that must be netted out before it means anything. The practical discipline follows directly: to track the research-peptide market, weight social and commercial signals as the leading indicators and read the news as confirmation that a phenomenon has already matured enough to generate an event a newsroom recognizes. By the time the coverage arrives, the move it describes is usually years old.

17 · Advertising and commercialization

Advertising is the evidence stream this study is least able to measure directly, and the report says so plainly rather than overclaim. A manifest of collected peptide-related social advertisements is registered in the research infrastructure but has not been analyzed at the creative level for this edition, so what follows is characterization drawn from the adjacent data, the marketplace, the social corpus, the news composition, not a measured advertising audit. The magnitude of the spend is left deliberately unstated because the evidence to size it responsibly is not yet in hand.

What the surrounding evidence does show clearly is a **two-tier promotional market**, and the two tiers barely resemble each other.

At the **regulated tier**, the marketing is professional and increasingly mainstream. Telehealth clinics and compounding-pharmacy networks advertise GLP-1 therapy, and increasingly other peptides, through conventional digital channels: paid search, connected-TV, sponsored content, the ordinary machinery of consumer-health customer acquisition. The steady telehealth and compounding presence in the news corpus (roughly 2,200 and 2,600 articles respectively) is partly a reflection of that spend and the funding rounds behind it [news corpus]. This tier looks like health-tech marketing because that is what it is, and it operates under at least some advertising oversight.

At the **gray-market tier**, promotion looks completely different. It runs on search-engine optimization, affiliate content, coupon codes, influencer seeding, and community "sourcing" recommendations that deliberately blur the line between peer advice and paid placement. The social corpus is thick with the raw material of this tier: creators walking viewers through reconstitution and dosing, posts that carry a ritual "for research and educational purposes only" disclaimer while functioning as product demonstrations, and vendor-adjacent accounts that seed protocols. The two astroturfing cautions the price tracker attached to specific sellers are a small, concrete glimpse of the same dynamic from the supply side [RUO price tracker]: in a market where reviews and recommendations can be manufactured, the boundary between organic enthusiasm and advertising is genuinely hard to locate, and is sometimes designed to be.

The commercialization story, then, is one of **rapid, low-friction market formation running ahead of regulation.** Dozens of sellers, a dense affiliate and influencer layer, and a regulated telehealth channel expanding alongside the gray one, all of it assembled in a few years, mostly without the advertising oversight that governs approved pharmaceuticals, where every claim must be substantiated and every risk disclosed. The research-peptide tier makes efficacy-adjacent claims that no approved drug could make in an advertisement, precisely because its products are not drugs in the eyes of the regulator, which is the same legal posture that lets them be sold at all.

The investor and professional read. The two-tier structure is the strategic heart of the commercialization question. The regulated telehealth and compounding tier is where conventional capital can operate, where customer acquisition is measurable, and where the demand created by the gray market can be captured inside a compliant channel, which is why it is the more investable adjacency even though it addresses the same underlying appetite. The gray tier's promotional engine is potent and cheap, but it is built on the same legal contingency as the market it serves: its freedom to make claims depends entirely on the products remaining classified as research chemicals, and its reliance on influencer and affiliate seeding is exactly the surface an FTC or FDA enforcement action, or a single defamation or injury suit, would target first. Sizing this spend properly, using the registered ad manifest together with a paid ad-intelligence source, is the most valuable single extension of this section. For now the report characterizes the shape of the advertising market and declines to invent its size, which is the same discipline it has applied to every class of evidence in this part: report what the data can carry, and name the gap where it cannot.

PART V

People, culture and society

The preceding parts established a structural fact: the peptide story splits into a regulated pharmaceutical boom that the press covers and a research-chemical boom that lives on social media. This part turns from the structure to the people inside it. It asks who is actually reaching for these compounds, what they say they want, what they say happens, and how a set of claims about weight-loss drugs became a running commentary on modern life. The discipline of the whole report tightens here rather than loosens, because this is the evidence class where it is easiest to mistake volume for truth.

18 · Who is reaching for peptides, and what they report

For approved medicines, a researcher who wants to know who uses a drug and how it goes can consult prescription records, insurer claims, and adverse-event registries. For the research-peptide world none of that exists, and the closest available substitute is what people volunteer about themselves in public. The corpus contains **1,917 first-person experience reports**, machine-extracted from social posts, each tagged by sentiment, reported effect, and, where stated, dose and side effects [social listening]. These are anecdotes. They evidence what people *say*, which is a fact about discourse, not a fact about biology. Read with that constraint held firmly in place, they are the richest available portrait of the culture's motivations.

The first surprise is that the tone is not the uncritical hype a casual observer might expect. Across all 1,917 reports the sentiment split is **about 42% positive, 37% neutral, and 21% negative** (801 positive, 702 neutral, 403 negative, 11 mixed) [social listening]. A neutral-and-positive center of gravity with a real negative tail of one in five is the profile of a community swapping notes, not a fan club. People report disappointment, non-response, and unpleasant effects often enough that the aggregate reads as testimony rather than testimonial.

The compounds people report on are the same ones that dominate the social conversation, which confirms that the discourse and the self-reported *use* point at the same molecules: GLP-1 drugs, GHK-Cu, BPC-157, ipamorelin, semax, MOTS-c, TB-500, and CJC-1295 all sit near the top of the experience-report rankings [social listening]. The *content* of what people chase clusters into a small set of aspirations, and the cluster is a portrait of optimization rather than of illness:

- **Cognitive performance.** Focus, relief from "brain fog," easier task initiation, study and work stamina. This is the reported territory of semax and selank.
- **Recovery and repair.** Joints, tendons, injuries, post-training soreness, general "healing." The reported territory of BPC-157 and TB-500.

- **Body composition.** Fat loss, muscle retention, "vascularity," and the GLP-1 experience of reduced "food noise." The territory of the incretin drugs, the growth-hormone secretagogues, and MOTS-c.
- **Skin and hair.** Anti-aging, wound healing, regrowth, "glow." The reported territory of GHK-Cu and the cosmetic peptides.
- **Sleep and energy.** A cross-cutting theme that attaches to the growth-hormone axis and the bioregulators.

Representative reports, quoted verbatim from public posts, show the texture and the vocabulary. On GHK-Cu: *"I can see some baby hairs growing back which is progress."* On BPC-157 and TB-500 stacked for injury: *"cycling BPC157 and TB500 for pain management which has added 27.5kg per side to this incline bench press."* On semax: *"What people report is clarity, the mental static dropping, easier task initiation, less procrastination."* On the growth-hormone secretagogues: *"No DAC feels more natural, pairs really well with Ipamorelin, and gives you way more control."* [social listening]. The language is domestic and procedural, doses, cycles, stacks, results, the register of a hobby with a technical literature, not of patients describing a treatment.

The one-sided-enthusiasm tell, quantified

The base report flagged a pattern worth stating precisely because it inverts the naive reading of the data. When the experience reports are grouped into the *research peptides* the internet loves and the *approved, potent drugs* medicine actually prescribes, the two groups report their experiences very differently.

REPORT GROUP	EXPERIENCE REPORTS	POSITIVE	NEGATIVE	REPORTS LISTING A SIDE EFFECT
Research peptides (GHK-Cu, BPC-157, MOTS-c, TB-500, ipamorelin, CJC-1295, semax, GHK, epithalon)	279	39%	**10%**	**13%**
Approved potent drugs (semaglutide, tirzepatide, insulin, somatropin)	146	43%	**32%**	**29%**

Source: 1,917-record experience-report layer of the social corpus, grouped by compound. "Negative" is the share of reports the extractor tagged negative; the final column is the share of reports that populate a specific side-effect field. Anecdotal evidence class: this measures how people talk, not what the compounds do.

The gap is large and it points the opposite way from the intuition it invites. The unapproved research peptides draw **three times less** negative reporting than the approved drugs, and list a concrete side effect less than half as often. It would be easy to read that as "the research peptides are gentler and work better." The evidence supports the reverse inference. Powerful drugs with real pharmacology, semaglutide, insulin, growth hormone, produce effects large enough that people notice and report both the benefit and the cost. The GLP-1 reports are full of the downside in plain language: *"my face swells in the 1-5 days following injection"*; *"She was SICK, it was bad and I almost*

took her to the ER"; *"it felt like I was the definition of a slow responder"* [social listening]. Compounds with subtler, slower, or expectation-driven effects generate the frictionless, near-uniformly positive testimony that is the signature of novelty and belief. Side effects for the research peptides are not absent, nausea, fatigue, injection-site swelling, headaches, and one striking MOTS-c report of *"a strange fever... all my joints and my skin started hurting"* appear throughout, but they are outnumbered by enthusiasm in a way that should raise a careful reader's skepticism rather than lower it. A thin negative tail on an unproven compound is not evidence of safety. It is evidence that the effect is too small, too slow, or too suggestible to force an honest accounting.

> THE DISCIPLINE
>
> "People are reporting remarkable recovery from BPC-157" and "BPC-157 has been shown to speed recovery in humans" are different sentences, and only the first is supported here. A thousand positive posts about a compound with two registered human trials is a thousand data points about belief, not about biology. The experience layer is the best map we have of what the peptide culture wants and how it talks. It is not, and cannot be, evidence that any of it works.

19 · The subcultures reaching for peptides

"The peptide community" is a convenient phrase that hides an important truth: there is no single community. The corpus resolves into at least four overlapping subcultures, each with its own anchor compounds, its own preferred platform, and its own dialect. They borrow from one another, a bodybuilder microdoses a nootropic, a biohacker tries a cosmetic peptide, but their centers of gravity are distinct, and recognizing them is essential to reading the market. The platform data makes the division of labor visible: research peptides skew heavily toward Reddit, the corpus's deliberation venue (BPC-157 draws 141 Reddit posts against 36 on TikTok; ipamorelin 90 against 28), while the more aesthetic and performative compounds spread onto Instagram and TikTok [social listening].

Longevity and anti-aging medicine

The oldest and most institutionally serious of the four. Its adherents frame peptides as tools for extending healthspan, and its anchor compounds are the bioregulators (epithalon and the Khavinson family), MOTS-c, NAD+ precursors, and the growth-hormone secretagogues used for their reputed effects on sleep, recovery, and body composition. This subculture overlaps with a real clinical world, longevity clinics, functional-medicine practices, and the telehealth channel, and its vocabulary borrows the language of gerontology: "senescence," "mitochondrial function," "biological age." It is the subculture most likely to appear alongside physician supervision and lab testing, and the one the "Longevity" news topic (1,257 articles) most directly tracks [news corpus].

The fitness and bodybuilding world

The subculture with the deepest roots and the most technical folk-pharmacology. Peptides entered bodybuilding through the growth-hormone axis decades ago, and the community's expertise in injection, dosing, and cycling is inherited from that history. Its anchor compounds are BPC-157 and

TB-500 for recovery, the growth-hormone secretagogues (ipamorelin, CJC-1295) for lean-mass and sleep, and IGF-1. Its native platform is Reddit for the detailed protocols and Instagram for the physique documentation. Its dialect is the most established: compounds are "stacked" and "run," vials are "reconstituted," and the community's "research use only" framing surfaces in the recurring euphemism of dosing "my rat" rather than oneself, a linguistic tell that users know the products are not sold for human use and route around it in plain sight [social listening].

Aesthetic and skin medicine

The subculture with one foot in a regulated industry. GHK-Cu, a copper peptide with a genuine and legal history as a cosmetic-serum ingredient, is its anchor, and its distinguishing move is the migration of a topical cosmetic into injected "research" use on the strength of online enthusiasm. Its reported goals are hair regrowth, wound healing, skin quality, and anti-aging, and its platform mix tilts toward the visual networks: GHK-Cu is the most evenly spread compound in the corpus (116 Reddit, 93 TikTok, 90 Instagram), consistent with a compound that lives simultaneously in a Reddit protocol thread and an Instagram before-and-after [social listening]. This is the subculture where the line between a lawful cosmetic purchase and an unapproved systemic use is blurriest.

Biohacking and nootropics

The subculture most oriented to the mind and most entangled with the broader "quantified self" movement. Its anchor compounds are the Russian-developed cognitive peptides semax and selank, alongside a long tail of experimental molecules. Its reported goals are focus, mood, anxiety reduction, and cognitive stamina, and its dialect imports the nootropics scene's habits of self-experimentation and dose-logging. Semax's platform spread is telling, unusually strong on TikTok relative to the other research peptides (71 Reddit, 44 TikTok, 44 Instagram), reflecting a compound explained and performed to a general audience as much as debated among specialists [social listening]. This is the frontier where the peptide movement meets the wider self-optimization internet, and it operates on almost no Western clinical footprint at all (Part on the science established zero U.S.-registry trials for semax and selank).

> **WHY THE SEGMENTATION MATTERS COMMERCIALLY**
>
> These are not one market but four demand curves with different anchor compounds, price sensitivities, and channel preferences. A longevity clinic patient and a bodybuilding-forum sourcer buy different molecules, tolerate different prices, and respond to different regulatory events. Any commercial or policy read that treats "peptide users" as a monolith will misjudge both the size and the resilience of each segment.

20 · The GLP-1 spillover into everyday life

If the research-peptide subcultures are a story about a few hundred thousand enthusiasts, the GLP-1 spillover is a story the whole culture is telling about itself. Over the past two years a genre of claim has taken hold in which the weight-loss drugs are reshaping not just bodies but the economy: curbing drinking and addiction, emptying bariatric-surgery schedules, changing the grocery cart

and the restaurant check, altering what people wear, where they travel, even what airlines spend on fuel. The corpus contains a dedicated analytical layer built to measure this "Ozempic economy" as discourse, **8,414 tagged observations** of GLP-1 secondary and tertiary effects, sorted into twelve domains of everyday life [news corpus] [social listening].

REPORTED SPILLOVER DOMAIN	TAGGED OBSERVATIONS	WRITTEN IN CAUSAL LANGUAGE	BACKED BY ESTABLISHED CAUSAL EVIDENCE
Alcohol and drinking	1,477	266	**0**
Addiction and compulsion	1,198	209	**0**
Fitness and exercise	1,060	177	**0**
Bariatric surgery	998	246	**0**
Aesthetics ("Ozempic face")	947	245	**0**
Attitudes to obesity	559	119	**0**
Clothing and apparel	536	86	**0**
Travel	525	78	**0**
Employment and work	371	69	**0**
Grocery spending	309	57	**0**
Restaurants	281	40	**0**
Consumer spending	153	19	**0**

Source: GLP-1 effect-observation layer of the corpus, 8,414 observations. "Causal language" flags text that asserts the drug caused the effect; "established causal evidence" flags an observation backed by a causal study. Every value in the final column is a database fact, not a rounding: the field is uniformly false across all 8,414 rows. Each observation is a media or social claim, never a transactional or clinical measurement.

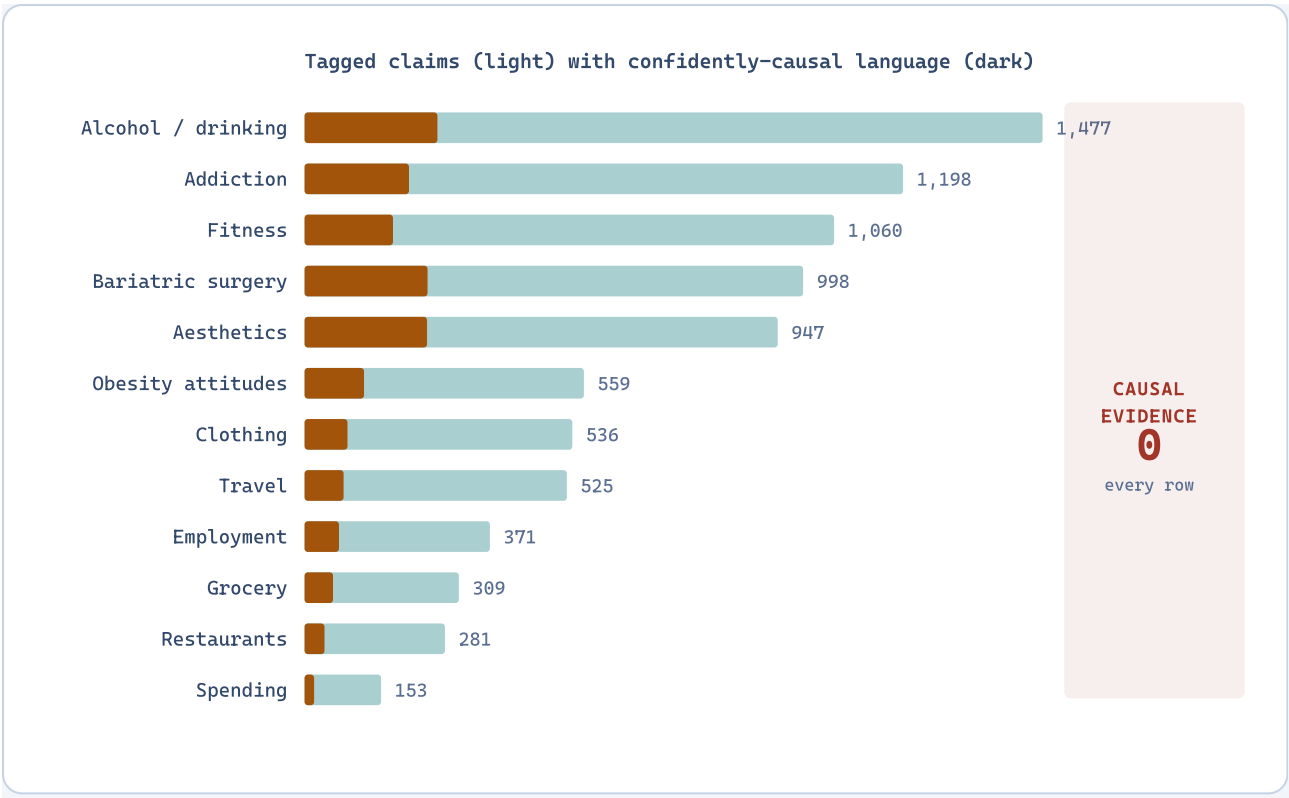

FIGURE 15 A conversation running ahead of its evidence. Tagged media and social claims that GLP-1 drugs affect each life domain, with the portion written in confidently causal language shaded darker. Across all 8,400+ observations, the count backed by established causal evidence is zero. Source: study corpus.

The final column is the whole point. Across every one of the twelve domains, and across all 8,414 observations, the number of claims backed by established causal evidence is **zero**. A meaningful minority, roughly one claim in six, is written in confidently causal grammar (*"Ozempic is killing the alcohol industry," "these drugs are ending food cravings"*), but not one of those assertions, as captured here, rests on a demonstrated causal study. What the corpus documents with precision is a vast, energetic *conversation* about the drugs' ripple effects, much of it phrased in the language of settled fact, resting on almost none of the evidence that language implies.

This is a finding about discourse, and it needs to be handled in both directions at once. It does not mean the effects are unreal. Several are biologically plausible and under genuine study; reduced alcohol craving, for instance, is an active research question with early supportive signals, and a drug that suppresses appetite will mechanically change some grocery and restaurant behavior at the margin. The claim here is narrower and more disciplined: **the cultural certainty has run years ahead of the causal evidence.** The "Ozempic changes everything" narrative is itself a major cultural phenomenon, worth studying precisely *as* a phenomenon, an early, unusually vivid case of a society metabolizing a new technology in real time. But it should be read as a population talking excitedly about a set of possibilities, not as a settled account of what these drugs are doing to the world.

For a professional the distinction is not pedantry, it is the entire risk. Several of the "Ozempic economy" claims have already moved markets: packaged-food and restaurant equities have been re-rated on the theory that GLP-1 users eat less, and medical-device and bariatric franchises on the theory that the drugs substitute for surgery. Those are multi-billion-dollar bets placed, so far, on a body of evidence whose causal column reads zero. When the controlled studies arrive, some of these

narratives will be confirmed, some will be quietly abandoned, and the gap between the two is where both the opportunity and the mispricing live. Section coverage of leading indicators in Part VI returns to exactly this point.

PART VI

The investor and professional lens

The first five parts described a phenomenon. This part reads it as a market. It maps where the money sits along the peptide value chain, asks honestly what the evidence can and cannot say about size, traces the competitive dynamics on both sides of the regulated line, lays out the regulatory risk that hangs over each layer, and closes with the specific signals a professional should watch. The register shifts toward analysis, but the evidence discipline does not. Where a number is unknowable from this data, the section says so rather than manufacturing one, and no claim here should be read as investment advice.

21 · The peptide value chain: where the money is

A useful way to see the industry is as a stack of layers, each performing a distinct function, capturing a distinct margin, and carrying a distinct risk. The corpus and the price tracker illuminate some layers brightly and others only at the edges, and the map below marks which is which.

1. Pharmaceutical innovators. The top of the regulated stack: Novo Nordisk and Eli Lilly, with Boehringer Ingelheim, Amgen, and Pfizer behind them. They own the molecules that matter commercially, the incretin drugs, and the evidence base that justifies them. This layer is where nearly all the clinical capital sits: semaglutide alone carries 741 registered trials and liraglutide 513, and the nine GLP-1 and incretin compounds together account for roughly 2,196 trials [ClinicalTrials.gov]. It is also where nearly all the scientific capital sits, about 20,000 PubMed publications for the GLP-1 class, some 2,500 in 2024 alone [PubMed]. This is the layer with a real moat: patents, trial data, manufacturing scale, and regulatory approval that no gray-market vendor can replicate.

2. Contract manufacturers and API suppliers. The least visible layer and one of the most important. Approved-drug innovators run their own or contracted GMP manufacturing; the entire gray market, by contrast, is fed by bulk active-pharmaceutical-ingredient suppliers, disproportionately in China. The price tracker's wholesale layer, though small (334 observations across six named bulk suppliers), captures the shape of it directly: the sourced supply splits almost evenly between **United States (180 observations) and China (154)**, and bulk powder is priced in

dollars per *gram* where retail vials are priced in dollars per *milligram* [RUO price tracker]. One catalogued Chinese supplier lists 1,500 mg of peptide at $9.50 at a minimum order of ten units; the same compound retails, reconstituted and vialed, at several dollars per milligram. The structural markup from bulk API to finished research vial is enormous, and it is the clearest indication of where the gray market's margin actually lives: not in manufacturing, but in vialing, branding, and distribution.

3. Compounding pharmacies. The regulated channel that the GLP-1 shortage pulled into the center of the story. When branded semaglutide and tirzepatide went into shortage in 2023 to 2025, U.S. compounding pharmacies (503A and larger 503B outsourcing facilities) were permitted to prepare copies, and an entire prescribing economy grew up around that allowance. The corpus tracks the footprint: **"Compounding" attaches to 2,566 articles** [news corpus], a steady beat of coverage that is also, increasingly, a litigation and enforcement beat as the shortages resolve and the legal basis for compounding narrows. This layer's economics are entirely contingent on regulatory status, which makes it the most exposed of the regulated tiers.

4. Telehealth platforms. The customer-facing front end of the regulated channel, and the fastest-growing consumer touchpoint. Online clinics turned peptide prescribing into a subscription funnel, pairing a brief clinical consult with fulfillment, and the corpus registers the scale of the phenomenon: **"Telehealth" attaches to 2,214 articles** [news corpus]. This is the layer where GLP-1 demand is monetized at the highest margin per patient, and the one most likely to broaden from incretins into other peptides as clinics look for the next subscription category. It is the regulated, visible cousin of the gray market, and its growth is the closest thing this study has to a hard signal that supervised peptide use is expanding (Part on real-world use).

5. Research-use-only and gray-market vendors. The layer this study observes most directly. The tracker has catalogued more than **1,000 candidate seller domains**, verified **443** as retailers, and captured live pricing from **302**, spanning roughly **18,400 listings across 114 compounds** [RUO price tracker]. The storefronts are commodity e-commerce, standard WooCommerce and Shopify installations rather than anything bespoke, which is itself a signal of how low the barrier to entry has become. This is a broad, shallow, low-concentration layer with essentially no moat, examined in detail below.

6. Testing and certificate-of-analysis labs. The trust infrastructure the gray market lacks. In a market of unapproved products with no regulatory identity guarantee, third-party purity and identity testing is the only substitute for a label you can trust, and the corpus shows how thin it is. Of the retailers the tracker checked, the large majority returned **no meaningful community-sentiment signal**, a handful showed mixed or negative sentiment, and two carried an explicit **suspected-astroturfing** caution [RUO price tracker]. Provenance claims are similarly sparse: fewer than a dozen of more than a thousand catalogued sellers make an explicit U.S.-manufacture claim on their homepage, and the field is blank for the overwhelming majority. A market this opaque is one in which independent testing is both the scarcest resource and, potentially, the most defensible business to build.

7. Ancillary supply. The picks-and-shovels layer: bacteriostatic water, syringes and sharps, alcohol swabs, reconstitution calculators, dosing apps, and the affiliate-content and coupon ecosystem that monetizes the discourse. Individually trivial, collectively the connective tissue that makes at-home use possible, and the layer least exposed to any single regulatory action because none of its components is itself a drug.

WHERE THE MOATS ARE, AND ARE NOT

Reading the stack top to bottom, defensibility collapses as you descend. The innovator layer is protected by patents, trial data, and approval. The compounding and telehealth layers are protected only by a regulatory permission that can be withdrawn. The RUO vendor layer has no moat at all, only price. The two layers with the most durable independent-business logic are the ones the gray market conspicuously lacks: trusted testing and trusted fulfillment.

22 · Market structure and sizing: what the evidence can and cannot support

A professional's first question is how big the market is. The honest answer for the research-peptide segment is that **no audited market-size figure is derivable from this evidence base, and this report will not invent one.** There is no prescription database, no sales-tax series, no trade association, and no public filings for a market that is, by design, informal. Anyone quoting a dollar TAM for research peptides is extrapolating from assumptions, not measuring. What the evidence *can* support is a rigorous read of *relative* scale and market *structure*, which is more useful for most decisions than a fabricated total would be.

On relative scale, the corpus weights are unambiguous. Measured by any regulated-market proxy, the GLP-1 class dwarfs everything else: it holds roughly 93% of compound-linked news coverage, the overwhelming majority of trials and publications, and the attention of the financial press. The research-peptide segment is, in regulated-market terms, small. Measured by *culture and commerce*, the ranking inverts: non-GLP-1 peptides hold roughly 80% of compound-linked social posts, and in the commercial market-strength model unapproved research chemicals (MOTS-c, BPC-157, GHK-Cu) rank alongside or above semaglutide [social listening] [P08 market model]. The two segments are not competing for the same dollar; they are different businesses at different scales, and pooling them produces a number that describes neither.

On structure, the tracker supports specific, defensible statements about the RUO layer:

STRUCTURAL MEASURE	VALUE	WHAT IT INDICATES
Priced sellers captured	302 (of 443 verified, 1,000+ candidate domains)	A broad, still-expanding vendor base; the figure is a floor
Listings / distinct compounds	~18,400 / 114	Deep catalogs, wide compound coverage
Sellers stocking the most-covered compound	44 (BPC-157)	Intense competition on flagship molecules
Top-three-seller share of "cheapest-price" wins	~53%	Moderate concentration on price, no dominant brand

STRUCTURAL MEASURE	VALUE	WHAT IT INDICATES
Storefront platforms	WooCommerce and Shopify (commodity)	Very low barrier to entry

Source: RUO price tracker and companion market-structure model, collection window July to August 2026. Counts exclude bot-walled and login-gated sellers, so every figure is a lower bound on the true market. These are commercial-availability measures, not usage or revenue.

The picture is a young, fragmented, low-concentration commodity market. Dozens of sellers stock the same flagship compounds on identical off-the-shelf storefronts, price is the primary axis of competition, and no brand has established a durable lead. That is not a market with a moat; it is a market in the land-grab phase, where entry is cheap, differentiation is thin, and the eventual consolidators, if consolidation comes, will most plausibly win on trust and fulfillment rather than on selection or price.

23 · Competitive dynamics on both sides of the line

The peptide industry runs two competitions at once, and they could hardly be more different.

The regulated side is a two-horse arms race. The incretin market is effectively a duopoly between Novo Nordisk (semaglutide, liraglutide, and the amylin combination CagriSema) and Eli Lilly (tirzepatide, plus the two candidates most likely to define the next round), with a credible third entrant in Boehringer Ingelheim (survodutide) and others behind. The competitive action is visible in the trial registry: Lilly's triple-agonist retatrutide (33 registered trials, all since 2019) and its oral candidate orforglipron (50 trials) are the pipeline threats to Novo's franchise, while CagriSema is Novo's answer [ClinicalTrials.gov]. This is classic pharmaceutical competition, fought with capital, trial data, and manufacturing scale, and the barriers are so high that only a handful of firms can play. An oral GLP-1 that reached approval would be the single most disruptive event, removing the injection barrier that has shaped the entire category and widening the addressable market again.

The gray side is the opposite: perfect competition with no barriers. As the structure data showed, the RUO market is fragmented across dozens of interchangeable vendors selling the same molecules from the same storefronts at dispersed prices. Even within a single vial size the spread is severalfold, BPC-157 at 10 mg ranges from roughly $2.30 to well above $20 per milligram across nearly 200 in-stock observations, clustered around a $7.50 median [RUO price tracker]. Wide, persistent price dispersion for an identical product is the textbook signature of a market with many entrants, weak brand power, and no pricing moat. Nobody in this layer has durable share; the "leaders" are simply whoever is cheapest this month.

The two sides meet at the channel line. The most interesting competitive dynamic is not within either market but *between* them: the regulated compounding-and-telehealth channel and the RUO gray channel compete for the same end user, on opposite terms. The regulated channel offers a prescription, physician oversight, and a defensible supply, at higher cost and under a regulatory permission that is narrowing. The gray channel offers price, selection, and no gatekeeper, at the cost of unverified identity and purity and real legal ambiguity. Where the shortage-era compounding

allowance ends, some demand will migrate back to branded product, and some will migrate *out* to the gray market. Which way it flows is the single most consequential open question for every layer below the innovators.

24 · The regulatory risk map

Regulatory status is the master variable of this industry. It determines which layers exist, how large they can grow, and how fast they can be erased. The corpus shows the topic is already central, not peripheral: **"Regulatory" attaches to 14,633 articles, "Safety" to 13,302, and "Litigation" to 3,098** [news corpus]. Mapping the principal risks onto the value chain:

- **FDA action in the compounding context.** The most immediate and most concrete risk. The compounding allowance that built the telehealth-and-compounding economy was contingent on official drug shortages; as those resolve, the legal basis contracts, and enforcement follows. This risk falls almost entirely on layers 3 and 4 (compounding pharmacies and telehealth), and a decisive move there would reprice both overnight while barely touching the innovators, who benefit as demand returns to branded product.
- **The RUO "research use only" gray zone.** The legal fiction that holds the gray market together: compounds sold "for laboratory research, not for human consumption" occupy a space that is lawful to sell and unlawful to market for human use. It persists on the strength of that labeling and of limited enforcement attention. Any change, an FDA import alert, a high-profile safety event, a decision to treat the vendors as unapproved-drug sellers rather than chemical suppliers, would hit layer 5 hardest and cascade up to the API suppliers in layer 2.
- **Import and enforcement risk.** Because so much bulk API originates overseas, customs and import policy is a structural chokepoint. The China-heavy supplier base means a tightening of import enforcement could disrupt supply well upstream of any action against the retailers themselves.
- **Off-label and advertising enforcement.** The regulated channel's promotional practices, and the gray channel's affiliate-and-influencer seeding, both sit in contested advertising territory. The two astroturfing cautions the tracker attached to specific sellers are a small concrete glimpse of a promotional layer operating largely without the oversight that governs approved pharmaceuticals.

> THE ASYMMETRY THAT MATTERS
>
> A single regulatory decision does not fall evenly. Tightening compounding rules helps the innovators and hurts the telehealth and compounding layers. An RUO crackdown guts the gray vendors and their API suppliers while leaving the regulated stack intact. Nothing plausible on the horizon threatens the innovator layer, whose moat is the approval itself. In peptides, regulatory risk is concentrated in exactly the layers with the least defensibility, which is why the gray market's low barriers to entry are matched by an equally low barrier to its removal.

25 · Leading indicators a professional should watch

The most valuable output of this report for a professional is not a forecast but a set of instruments. The study's own structural findings imply which signals lead and which lag, and the central lesson, that culture and commerce showed up years before the news, means the conventional indicators (mainstream coverage, sell-side notes) are the *lagging* ones. The signals worth watching, roughly in the order they move:

- **Social diffusion ahead of news.** The report's clearest structural finding is that social attention leads mainstream coverage, by as much as five years for BPC-157. A compound's rising social share of voice, well before any newsroom notices, is the earliest available signal that a new molecule is entering the culture. Watch the social series, not the press.
- **RUO vendor and listing counts.** Vendors do not stock what no one buys. The number of sellers carrying a compound, and the direction of that count over time, is a real-time proxy for demand that requires no survey. A compound moving from a handful of sellers to dozens is a demand signal; a thinning vendor base is an early warning of either saturation or regulatory pressure.
- **Trial registrations, especially by new sponsors.** The single most important forward signal in the science is the arrival of a genuine efficacy trial where none existed. The concentration of nearly every *first* real interventional study of the trendy peptides in one sponsor (Hudson Biotech, all starting February 2026) is precisely the kind of event to watch: it marks the moment a research chemical begins the transition toward, or away from, legitimacy [ClinicalTrials.gov].
- **Telehealth and compounding coverage as a channel proxy.** The steady news beat around telehealth (2,214 articles) and compounding (2,566) tracks the growth of the regulated channel. Inflections there, a surge in enforcement litigation, a wave of clinic launches or closures, lead changes in where supervised demand flows.
- **The first genuine efficacy readouts.** The whole science-versus-culture gap collapses, compound by compound, the day a controlled human trial reports. For any research peptide, the readout of its first real efficacy study is the binary event that will either legitimize it or deflate it. Those readouts, beginning to register from 2026, are the events that will resolve the report's central tension.
- **Causal studies behind the "Ozempic economy" claims.** For the spillover narratives that have already moved consumer-sector equities, the leading indicator is the arrival of the first causal evidence in each domain. The corpus shows that column at zero today; the first well-designed study to move it, on alcohol, on grocery spend, on bariatric substitution, will convert a market-moving story into a market-moving fact, or refute it.

These indicators are analytical guidance derived from the report's own signals, not recommendations. "Watch" identifies where the diffusion and evidence dynamics are most active, and implies no view on the purchase, sale, or use of any compound, security, or company. This report is market, media, and cultural analysis and is not investment or medical advice.

PART VII

Science versus hype: the compound dossier

26 · The dashboard, and why there is no single score

Every reader who reaches this point wants the same thing: one number per compound, a leaderboard, a *Peptide Adoption Index* that says semaglutide is a 97 and epithalon is a 44 and settles the matter. This report refuses to build it, and the refusal is not modesty. It is the finding.

Adoption is not one quantity. A compound can carry deep science and almost no culture, or enormous culture and almost no science, and those two cases are the most interesting objects in the entire dataset. A composite would average them into a forgettable middle and destroy exactly the divergence worth seeing. Worse, the signals are measured in incommensurable units on evidence of wildly uneven quality: a PubMed count and a TikTok mention count and a market-strength score do not share a denominator, and any formula that added them would smuggle an editorial weighting in through the back door and present it as a measurement. So the report publishes the **six component signals side by side** and lets the reader watch where they agree and, more often, where they violently disagree.

Six questions, six signals, each answering something the others cannot:

- **Scientific activity.** Is it being studied at all? *(PubMed lifetime and 2024 publication counts [PubMed].)*
- **Clinical development.** Is it being tested in actual humans? *(Genuine compound-specific registered trials [ClinicalTrials.gov].)*
- **Commercial availability.** Can you buy it, and is the market deep? *(0 to 100 market-strength model [P08 market model]; live listing and seller counts [RUO price tracker].)*
- **Search and buying intent.** Are people trying to learn about it or purchase it? *(Search-intent proxy inside the market model [P08 market model].)*
- **Social attention.** Are people discussing it? *(Share-of-voice mention counts [social listening].)*
- **Media attention.** Is the press covering it? *(Share-of-voice mention counts [news corpus].)*

Laid out together for the sixteen compounds this dossier profiles, the columns pull in different directions, and the disagreement is the whole point.

COMPOUND	EVIDENCE STATUS	PUBMED (LIFE / 2024)	GENUINE TRIALS	MARKET STRENGTH	SOCIAL	NEWS
Semaglutide	Approved	5,465 / 1,052	741	90	286	17,186
Tirzepatide	Approved	2,391 / 399	274	94	244	12,291
Retatrutide	Investigational	169 / 45	33	98	125	2,153
Orforglipron	Investigational	127 / 17	50	not scored	68	2,786

COMPOUND	EVIDENCE STATUS	PUBMED (LIFE / 2024)	GENUINE TRIALS	MARKET STRENGTH	SOCIAL	NEWS
Cagrilintide	Investigational	97 / 14	43	77	56	572
BPC-157	Research chemical	228 / 7	2	93	274	1,098
TB-500	Research chemical	1,073 / 15	18	82	206	530
GHK-Cu	Research chemical	189 / 9	1	93	299	207
MOTS-c	Research chemical	250 / 36	1	93	179	269
Ipamorelin	Research chemical	54 / 2	2	82	164	377
CJC-1295	Research chemical	33 / 0	1	59	141	166
Tesamorelin	Approved (niche)	121 / 4	24	94	114	81
PT-141	Approved	121 / 8	10	73	~0	38
Epithalon	Research chemical	142 / 1	0	89	152	0
Semax	Research chemical	231 / 4	0	81	159	~219
5-Amino-1MQ	Research (non-peptide)	3 / 0	0	not scored	63	10

Sources: PubMed lifetime and 2024 publication counts [PubMed]; genuine compound-specific interventional trials, discounting synonym-collision matches [ClinicalTrials.gov]; 0 to 100 market-strength score [P08 market model]; social and news share-of-voice mention counts [social listening; news corpus]. TB-500 trial and publication counts belong to the thymosin beta-4 pharmaceutical program, distinct from the systemic fragment sold to consumers (see its entry). "not scored" means the compound is absent from the market model. Counts are a snapshot as of August 2026 and are not combined into any composite.

Run your eye down the columns and the structure jumps out. The *News* column is a GLP-1 monopoly: semaglutide and tirzepatide alone carry more than 29,000 mentions, and every four-figure entry is an incretin drug. The *Social* column is its near-inverse, topped by GHK-Cu, BPC-157, TB-500, and MOTS-c, none of them approved for the uses being discussed. The *Genuine trials*

column separates the approved world (hundreds of trials) from the research-chemical world (zero, one, or two), and it does not care in the slightest what the *Market strength* column says, which is the scandal the next sixteen entries document. A single index would have blended all of that into a shrug.

27 · The four quadrants, and the geometry beneath them

Collapse the six signals onto two composite axes, not to score compounds but to *place* them, and the field organizes itself. The first axis is **evidence**: how much validated science stands behind a compound, running from approved-and-trial-tested at one end to unproven research chemical at the other. It is built from the two signals that track truth rather than attention, clinical trials and scientific literature. The second axis is **momentum**: how much commercial and cultural energy a compound has gathered, built from market strength, search intent, and social attention. Plot every compound on those two axes and the plane divides into four quadrants.

- **Proven and popular** (high evidence, high momentum). The GLP-1 drugs. Real trials, real approval, and real attention arrived together and reinforce each other. This is adoption in its healthy form, where the science and the enthusiasm point at the same molecule.
- **Proven but quiet** (high evidence, low momentum). Older approved peptides such as tesamorelin and bremelanotide, and the broad class of hormone therapies. The evidence exists; the culture never formed. These are drugs, not phenomena.
- **Unproven but booming** (low evidence, high momentum). BPC-157, GHK-Cu, TB-500, MOTS-c, ipamorelin, CJC-1295, epithalon, semax, 5-Amino-1MQ. Intense online culture and deep commercial availability standing on a clinical base of one trial, or none. This is the hype quadrant, and it is by far the most crowded.
- **Neither** (low evidence, low momentum). The long tail of the 301 catalogued compounds that neither science nor the market has embraced.

> THE CORE DIAGNOSIS
>
> The peptide phenomenon is really the story of two quadrants. The GLP-1 drugs sit in "proven and popular." The research-peptide movement lives almost entirely in "unproven but booming," where availability and enthusiasm have far outrun the evidence. Naming a compound's quadrant tells you not just how much traction it has but what kind, and whether that traction rests on anything at all.

The dossiers that follow walk the compounds in roughly the order the market and the culture would rank them, GLP-1 leaders first, then the research-peptide core, then the approved outliers and the fringe. Each entry holds the same shape: what the compound is and how it works in plain language, what the evidence actually shows in numbers, how strong its market is, the shape of its diffusion across news and social, what the discourse sounds like, and the adoption and investor read. The goal is to make each compound legible on its own terms, because the most useful fact about any of them is usually not how much attention it has but the mismatch between its columns.

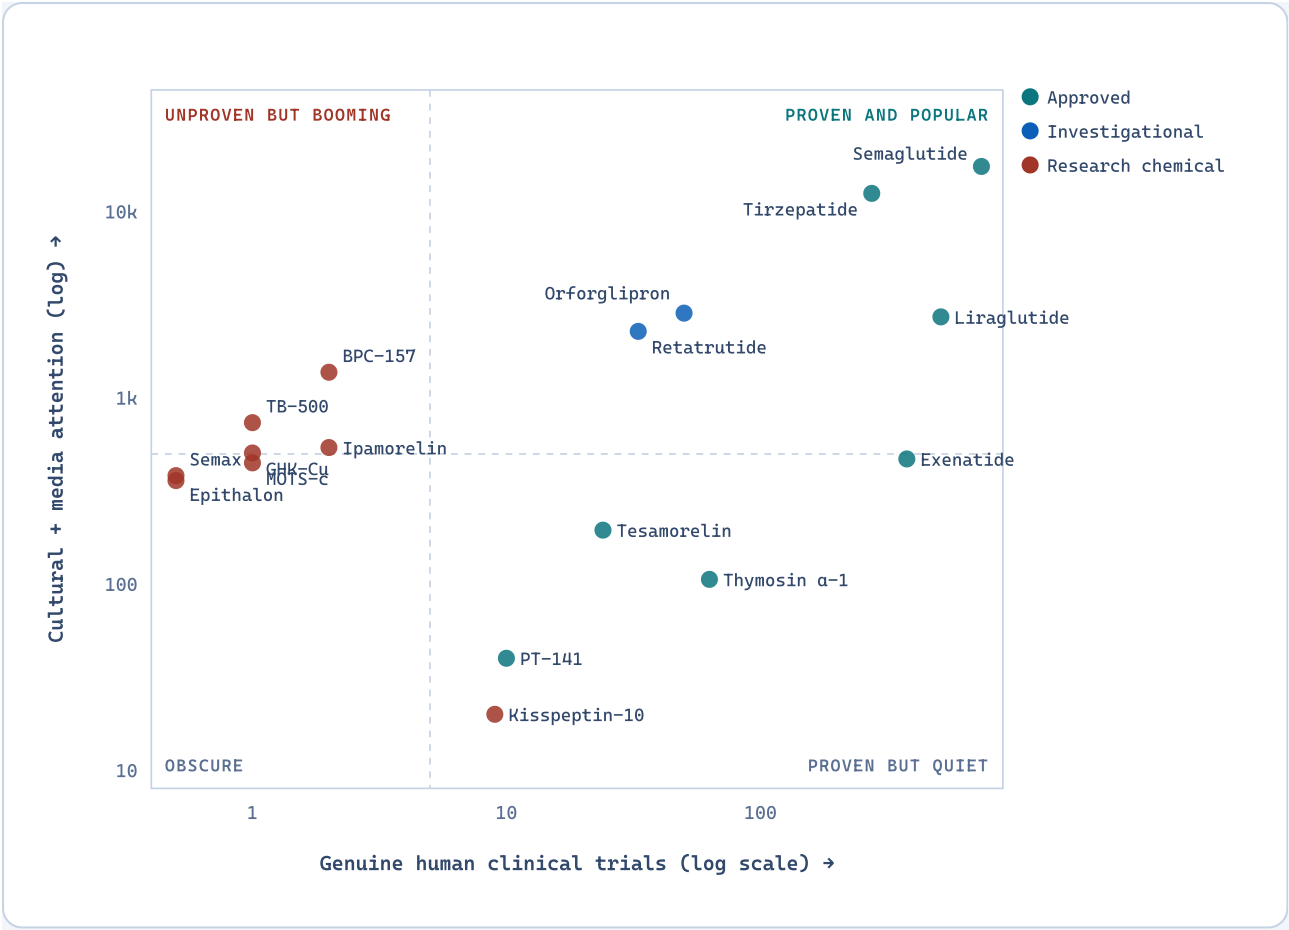

FIGURE 16 Science versus culture. Each compound placed by the number of genuine human clinical trials (horizontal, log scale) against its cultural and media attention (vertical, log scale). GLP-1 drugs sit top-right, proven and popular; the research peptides cluster top-left, unproven but booming. Trials: ClinicalTrials.gov. Attention: study corpus.

28 · Semaglutide

What it is. The active ingredient in Ozempic, Wegovy, and Rybelsus, and the molecule that made the word "peptide" a household term. Semaglutide is an engineered mimic of GLP-1, a gut hormone released after eating that prompts insulin release, slows stomach emptying, and signals fullness to the brain. The engineering trick is durability: natural GLP-1 lasts minutes, semaglutide lasts about a week, which is what turns a fleeting satiety signal into a standing weight-loss drug.

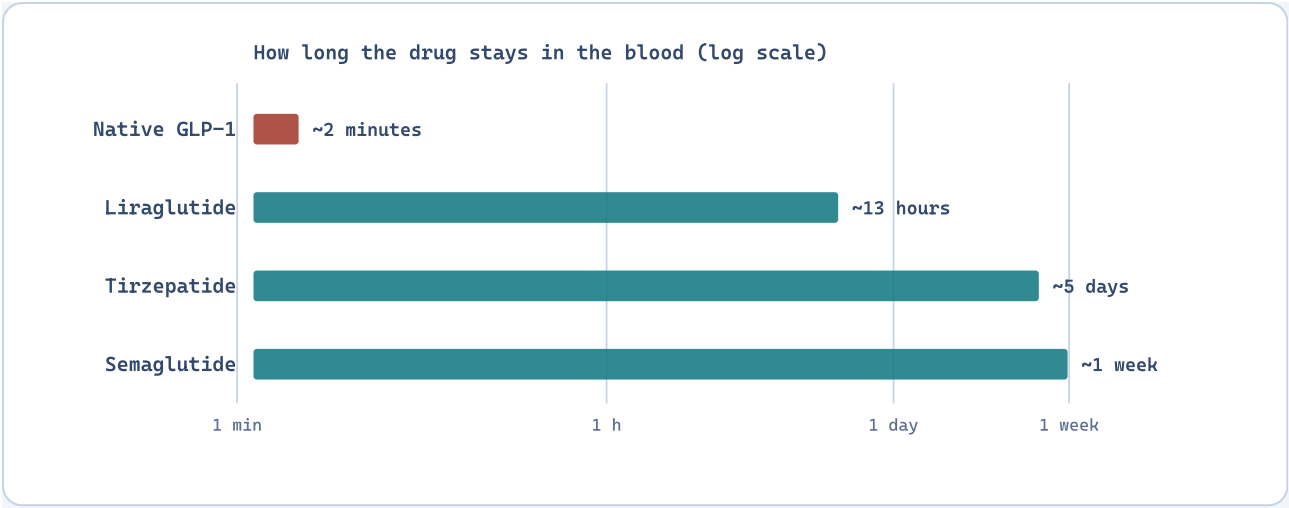

FIGURE 17 Engineering the clock. Circulating half-life on a log scale: natural GLP-1 lasts about two minutes, while the engineered drugs, shielded by an added fatty-acid chain, last from roughly half a day to a week, which is what makes a weekly injection possible. Source: SBL incretin monographs.

Evidence. This is the most heavily evidenced compound in the entire study. ClinicalTrials.gov lists **741 registered trials**, 332 of them completed, spanning Phase 1 through Phase 4 with start years from 2007 to 2028 [ClinicalTrials.gov]. PubMed indexes **5,465 publications, 1,052 of them in 2024 alone** [PubMed]. The literature was near-flat before 2018 and has gone nearly vertical since, the fingerprint of a genuine, well-funded scientific boom. The efficacy behind that boom is concrete: in the pivotal STEP-1 obesity trial, semaglutide produced a mean **14.9% body-weight reduction** over 68 weeks against 2.4% on placebo (Wilding et al., 2021), and in the SELECT trial it cut major cardiovascular events by about **20%**, the result that earned its March-2024 heart-disease indication and changed what kind of drug it is.

Market. A market-strength score of **90** and roughly **735 live research-use-only listings across 167 sellers** at a median near $16.83 per milligram [P08 market model; RUO price tracker]. Notably, semaglutide is *not* the top of the RUO market, several research chemicals outrank it, because the gray market prices attention and availability, not the FDA label.

Diffusion. Textbook top-down. It surfaces in both news and social from mid-2019, but the news volume dwarfs the social by roughly sixty to one (17,186 news mentions against 286 social) [news corpus; social listening]. The public met this drug through business and health coverage, not through a subculture.

Discourse. Among first-person experience reports the sentiment is strikingly *balanced*: of 37 tagged reports, 18 positive, 6 neutral, and 13 negative [social listening]. That balance is a mark of a genuinely potent drug, people report the dramatic appetite suppression and the nausea, vomiting, and "Ozempic face" with equal candor, which is exactly what honest reporting of a real pharmacological effect looks like.

The read. Proven and popular, the anchor of that quadrant. The trajectory is institutionalization: broadening indications, insurance fights, and supply as the binding constraint rather than demand. For an investor, semaglutide is no longer the frontier; it is the incumbent that the rest of the pipeline is measured against.

29 · Tirzepatide

What it is. The active ingredient in Mounjaro and Zepbound, and semaglutide's more powerful successor. Tirzepatide is a *dual agonist*: it mimics GLP-1 and a second gut hormone called GIP simultaneously. In plain terms, it pulls two metabolic levers at once, which in head-to-head trials has translated into greater average weight loss than semaglutide.

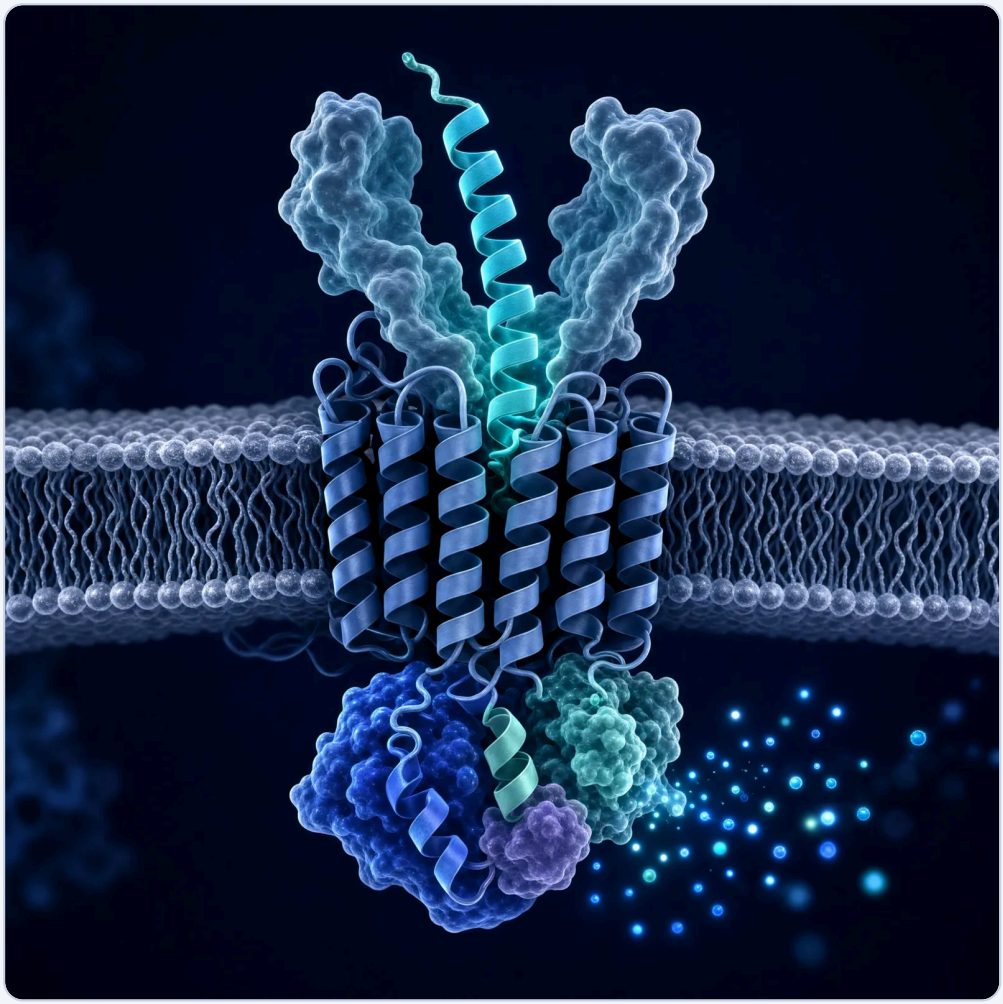

FIGURE 18 A hormone meeting its receptor. GLP-1 and its drug mimics act on a G-protein-coupled receptor threaded through the cell membrane; binding on the outside triggers a signaling cascade inside the cell. Illustrative schematic of receptor signaling, not a literal molecular structure.

Evidence. 274 registered trials, weighted toward the late-stage end that matters (50 Phase 3, 45 Phase 4), with rapid recent growth [ClinicalTrials.gov]. PubMed shows **2,391 publications, 399 in 2024**, up from just 11 in 2020, one of the steepest publication curves in the dataset [PubMed]. In its pivotal SURMOUNT-1 obesity trial, the 15 mg dose produced a mean **20.9% weight reduction** at 72 weeks (Jastreboff et al., 2022), the largest yet reported for an approved medicine and the reason a single trial readout can move a share price.

Market. The highest market strength of any *approved* drug here at **94**, with about **594 listings across 133 sellers** at the richest median in the set, roughly $22 per milligram [P08 market model; RUO price tracker]. The high per-milligram price reflects both potency (doses are small) and demand.

Diffusion. Top-down, slightly steeper than semaglutide. First news mid-2021, first meaningful social late 2022; 12,291 news mentions against 244 social [news corpus; social listening]. Again the newspaper leads the feed.

Discourse. Balanced, like its sibling: of 40 experience reports, 16 positive, 13 neutral, 11 negative [social listening]. The negatives are the familiar gastrointestinal complaints, reported without hesitation.

The read. Proven and popular, and currently the efficacy leader of the class. The trajectory question is not whether tirzepatide works but how it splits the market with next-generation triple agonists and oral options. Watch it as the benchmark every challenger, approved or gray-market, is priced against.

30 · Retatrutide

What it is. Eli Lilly's investigational "triple agonist," which adds a third lever, glucagon-receptor activity, to the GLP-1 and GIP mechanisms of tirzepatide. The glucagon arm is thought to raise energy expenditure, not just suppress appetite. In early trials it has produced the largest weight reductions yet reported for the class. It is **not approved for anyone**.

Evidence. Thin but fast-moving and legitimate: **33 registered trials**, 21 completed, evenly split between early and Phase 3 stages, all since 2019 [ClinicalTrials.gov]. PubMed holds **169 publications, 45 of them in 2024** from a base of zero in 2020 [PubMed]. This is a real, industry-run program in mid-flight, not a research-chemical ghost. In its phase-2 obesity trial the weight-loss curve had **still not flattened at 48 weeks**, reaching toward the range usually seen only with bariatric surgery, the deepest drug-induced weight loss reported to date.

Market. Here is the anomaly that defines it. Retatrutide tops the entire commercial model at **98 out of 100**, the single highest market-strength score in the study, and it carries the highest search-intent signal of any compound, roughly **2,080** [P08 market model]. The RUO market already lists it across about **728 products from 188 sellers** [RUO price tracker]. A drug approved for no one is the most commercially energized molecule in the field.

Diffusion. Top-down but running *ahead of approval*. First news late 2022, driven by trial readouts and investor anticipation; social follows in 2024 [news corpus; social listening]. The gray market did not wait for the FDA.

Discourse. Balanced and physically vivid. Of 25 reports, 7 positive, 11 neutral, 7 negative [social listening]. Users describe real potency and real side effects, including facial swelling and near-fainting hunger crashes: one wrote of feeling *"like I haven't eaten for a very long time and I'm getting close to fainting because of it"* [social listening]. This is the honest-reporting signature of an active drug, appearing years before the drug is legally available.

The read. The pre-approval supernova. It sits between "investigational" and "proven and popular," pulled toward the latter by trial momentum. It is also the cleanest case study in how the RUO market front-runs the pharmaceutical pipeline, selling a blockbuster-in-waiting before its label exists. Watch the Phase 3 readouts; watch, too, whether gray-market demand for an unapproved triple agonist draws regulatory attention.

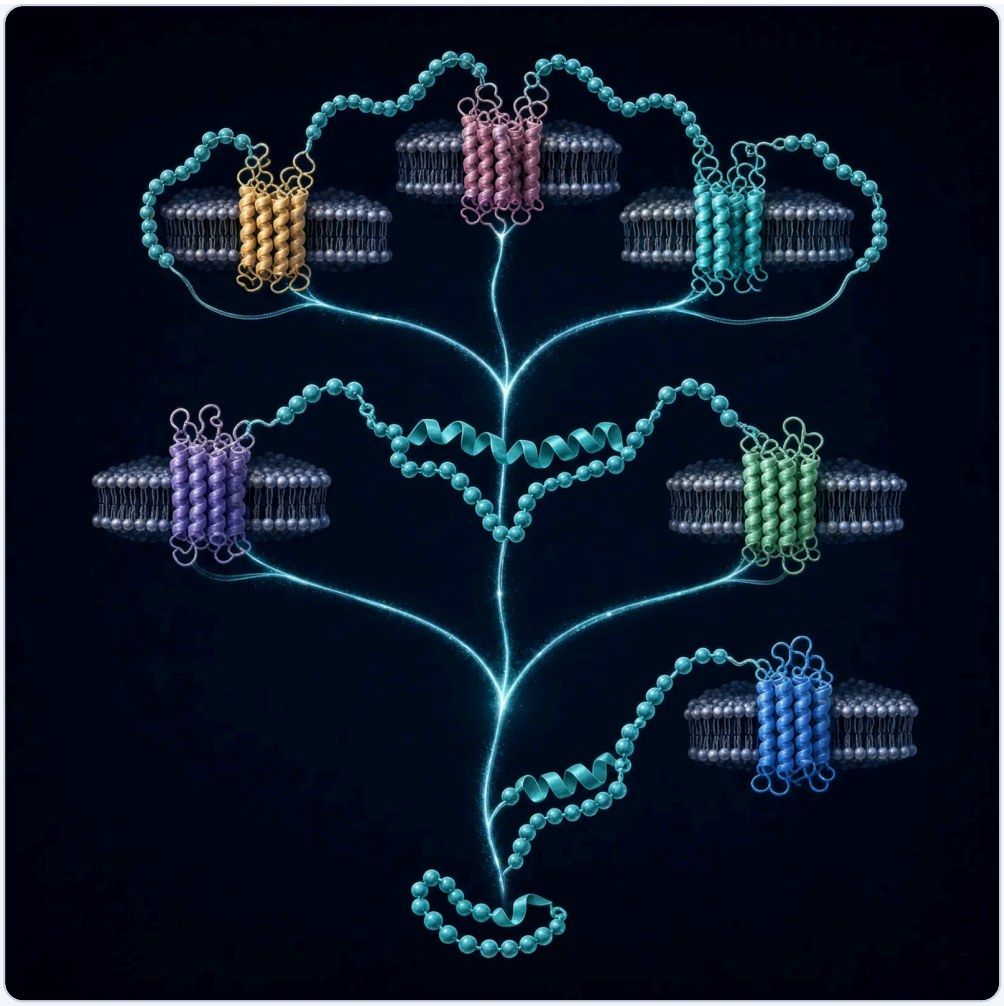

FIGURE 19 One hormone, many descendants. The incretin drugs are engineered relatives of a single natural signal, evolving from single-target (mono-agonist) molecules to dual- and triple-agonists that engage more receptors at once, which in trials has deepened the metabolic effect. Illustrative concept map, not a phylogeny or a set of literal structures.

31 · Orforglipron

What it is. An *oral, non-peptide* GLP-1 agonist, a small molecule that reproduces the GLP-1 effect in a daily pill rather than a weekly injection. If it reaches approval, it removes the needle from the equation, the single biggest friction point in the entire GLP-1 experience.

Evidence. A serious late-stage program: **50 registered trials, 25 of them Phase 3**, all since 2019 [ClinicalTrials.gov], with an early literature of **127 publications, 17 in 2024** [PubMed]. The science is real and advancing; the compound is investigational. Its significance is the delivery route: because it is a small molecule rather than a peptide, it survives the gut, where a swallowed peptide loses more than 99% of its dose (oral-peptide bioavailability is 0.4 to 1%), so it can work as a daily tablet. Its ATTAIN (obesity) and ACHIEVE (type-2 diabetes) phase-3 trials are reading out now.

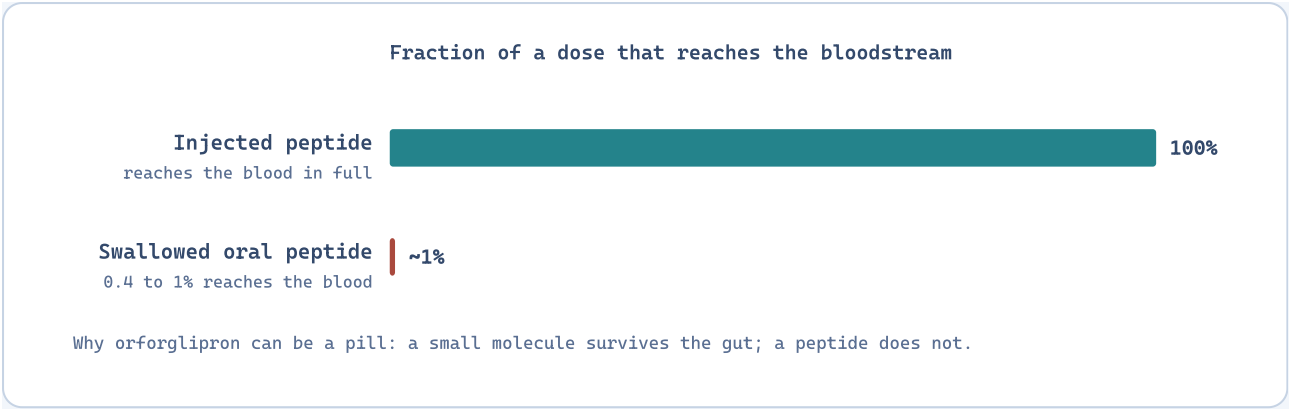

FIGURE 20 Why the pill matters. A swallowed peptide is almost entirely destroyed in the gut (only 0.4 to 1% of an oral-peptide dose reaches the blood), which is why peptide drugs are injected. Orforglipron is a small molecule that survives the gut and can be a tablet. Source: SBL orforglipron monograph.

Market. This is the most telling entry in the dossier. Orforglipron is **absent from the market-strength model entirely** and appears in the RUO price corpus as only **about 16 listings from roughly a dozen sellers** [P08 market model; RUO price tracker]. The gray market, which stocks unapproved injectables across hundreds of listings each, has barely touched it. An oral small molecule is harder to counterfeit and reconstitute than an injectable peptide, and it holds no appeal for a subculture built around vials and needles.

Diffusion. Top-down and news-only. It draws **2,786 news mentions**, more than BPC-157, against just **68 social mentions** [news corpus; social listening], the exact mirror image of the research peptides.

Discourse. Sparse. There is little first-person culture here, because there is little to reconstitute or "stack." The conversation lives in equity research, not on TikTok.

The read. A top-down pharmaceutical event, not a cultural one. It occupies the investigational edge of "proven and popular" and, if approved, is the compound most likely to *widen* that quadrant by pulling in the needle-averse majority. For investors it is the category-breaker to watch; for the peptide subculture it is a non-event, which is itself the point.

32 · Cagrilintide

What it is. A long-acting *amylin* analog. Amylin is a hormone co-secreted with insulin that promotes satiety through a different pathway than GLP-1. Cagrilintide is most advanced as half of **CagriSema**, a combination with semaglutide designed to attack appetite on two hormonal fronts at once.

Evidence. 43 registered trials, 21 of them Phase 3, though most are CagriSema combination studies rather than cagrilintide alone [ClinicalTrials.gov]. PubMed shows **97 publications, 14 in 2024** [PubMed]. A genuine, if younger, development program riding semaglutide's coattails.

Market. A market-strength score of **77**, with about **268 listings across 116 sellers** at a rich median near $13.64 per milligram [P08 market model; RUO price tracker]. Solid but a clear tier below the headline incretins.

Diffusion. Top-down, and very recent: first news mid-2024, first social mid-2025 [news corpus; social listening]. Combined with the separately tracked "CagriSema" name (1,324 news mentions), the franchise is a meaningful news presence but a minor social one.

Discourse. Limited first-person culture; the interest is largely anticipatory, tied to trial results and the amylin story as the "next mechanism."

The read. Investigational, trending toward proven-and-popular if the combination clears. It matters less as a standalone consumer phenomenon than as evidence that the obesity-drug race is broadening from single to multi-hormone mechanisms. Watch it as a pipeline signal rather than a culture signal.

33 · BPC-157

What it is. "Body Protection Compound 157," a synthetic chain of fifteen amino acids derived from a protein found in human gastric juice. It is marketed for healing, tendon and joint recovery, and gut repair. It is a research chemical: legal to sell for laboratory use, not approved for humans anywhere.

Evidence. The signature mismatch of the whole report. Despite total online ubiquity, BPC-157 has **two genuine human trials** (a 2015 Phase 1 and a Phase 2 that began recruiting in 2026) and a lifetime literature of just **228 publications, only 7 in 2024** [ClinicalTrials.gov; PubMed]. Almost all of that literature is animal and cell work. The reputation rests on rat studies and anecdote, not human evidence.

Market. Enormous. A market-strength score of **93**, and the **deepest commercial footprint of any compound in the study: 1,052 live listings across 240 sellers** [P08 market model; RUO price tracker]. More sellers stock BPC-157 than stock semaglutide. It is frequently sold in a blend with TB-500, which scores 91 on its own.

Diffusion. The definitive bottom-up case. BPC-157 appears in the social record from **November 2019** but is essentially absent from the news corpus until **mid-2024**, a lag of roughly five years between culture and coverage [social listening; news corpus]. The newsroom arrived half a decade after the forum.

Discourse. Enthusiastic and skewed positive: of 40 reports, 19 positive, 17 neutral, only 4 negative [social listening]. The texture is that of real use, with dosing protocols and before-and-after logs, one user reporting *"pain before 4 to 6 out of 10, pain now 2 to 3 out of 10, a 50 percent-plus reduction"* [social listening]. Compelling as testimony; not evidence of efficacy.

The read. The flagship of "unproven but booming," and the crossover to watch. It is the test case for whether online popularity alone can pull a compound toward formal scrutiny: the 2026 Phase 2 is the first real chance to find out whether the anecdotes survive a controlled trial. Still unapproved, still evidence-thin, but no longer invisible.

34 · TB-500

What it is. Sold as a systemic recovery and repair peptide, "TB-500" is a synthetic fragment (residues 17 to 23) of thymosin beta-4, a naturally occurring cell-repair protein. The consumer product and the pharmaceutical protein share a name and very little clinical overlap, a distinction that inflates its apparent evidence base.

Evidence. This is where the numbers must be read carefully. The **18 registered trials and 1,073 publications** attached to "thymosin beta-4" belong to a legitimate pharmaceutical program in eye drops, wound healing, and cardiac repair, run by companies like RegeneRx [ClinicalTrials.gov; PubMed]. Only *one* of those trials uses the "TB-500" fragment as sold to consumers, and it began in 2026. The systemic injectable that the fitness world buys has essentially **no completed efficacy trials**, and the underlying literature is *declining*, from 51 papers in 2015 to 15 in 2024 [PubMed].

Market. Market strength **82**, with about **515 listings across 212 sellers** at a median near $14.72 per milligram [P08 market model; RUO price tracker]. Its natural habitat is the BPC-157 stack.

Diffusion. Bottom-up. First social March 2021, first news mid-2024 [social listening; news corpus], a familiar multi-year lag, though shorter than BPC-157's.

Discourse. Positive-leaning and recovery-focused: of 23 reports, 10 positive, 11 neutral, 2 negative, with users crediting it for training recovery between sessions [social listening].

The read. Unproven but booming, and a cautionary tale in evidence laundering: the impressive-looking trial and publication counts describe a different molecule than the one being injected. Watch whether the 2026 fragment-specific trial closes that gap or widens it.

35 · GHK-Cu

What it is. A copper tripeptide, glycyl-histidyl-lysine bound to a copper ion, that occurs naturally in plasma and has a genuine, decades-long history as a *topical* skincare ingredient for wound healing and anti-aging. The new twist is a subculture injecting it systemically for the same claimed benefits.

Evidence. Modest and topical. **One genuine trial** (a topical Phase 2 starting in 2026) and **189 publications, 9 in 2024**, most concerning skin application rather than systemic injection [ClinicalTrials.gov; PubMed]. There is essentially no human evidence for the injected use that drives its online following.

Market. A market-strength score of **93** and search intent around **350**, among the highest of the research peptides, with roughly **722 listings across 247 sellers** at a low median near $3.28 per milligram, reflecting its cosmetic-commodity heritage [P08 market model; RUO price tracker].

Diffusion. Bottom-up with a cosmetic origin. First social July 2020, first news not until 2025; combined with the separately tracked "GHK" it is the **most-discussed non-GLP-1 compound in the entire social corpus** (299 mentions, plus 261 for GHK) against just 207 news mentions [social listening; news corpus].

Discourse. Overwhelmingly non-negative: of 51 reports, 21 positive, 28 neutral, only 2 negative [social listening]. The talk is skin, hair, and "glow," one user noting *"I can see some baby hairs growing back which is progress"* [social listening]. The near-total absence of negatives is itself a caution flag: subtle or expectation-driven effects generate frictionless testimony.

The read. The cosmetic crossover, and the clearest example of a compound migrating from an approved topical context into systemic "research" use purely on the strength of online enthusiasm. Watch it as the model for how a legitimate ingredient gets pulled across the needle by culture.

36 · MOTS-c

What it is. A mitochondrial-derived peptide, a sixteen-amino-acid chain encoded within mitochondrial DNA that appears to act as an "exercise mimetic," influencing metabolism and endurance. It is marketed for energy, fat loss, and athletic performance. Research chemical status.

Evidence. One genuine interventional trial (a Phase 2a beginning in 2026) among five nominal matches, most of which are observational biomarker studies that merely mention the peptide [ClinicalTrials.gov]. Its literature, at **250 publications and 36 in 2024**, is the fastest-growing of the niche peptides, though still a rounding error beside the incretins [PubMed].

Market. Market strength **93**, tied near the top of the research-chemical cohort, with about **495 listings across 217 sellers** and search intent near 280 [P08 market model; RUO price tracker].

Diffusion. Bottom-up. First social July 2020, first news 2025 [social listening; news corpus].

Discourse. Performance-driven and positive-leaning: of 29 reports, 14 positive, 12 neutral, 3 negative [social listening]. One endurance-focused user reported their *"run went through the roof, smashed 6 PBs"* [social listening], the kind of concrete performance claim that spreads fast in athletic communities.

The read. Unproven but booming, the metabolic-performance vanguard of the movement. The "exercise in a vial" framing gives it durable appeal. Watch the 2026 Phase 2a as one of the first real tests of whether any of the endurance claims hold.

37 · Ipamorelin

What it is. A growth-hormone secretagogue: a five-amino-acid peptide that mimics ghrelin to prompt the pituitary to release the body's own growth hormone. It is marketed for recovery, sleep, body composition, and anti-aging, usually stacked with a GHRH analog like CJC-1295.

Evidence. Two genuine Phase 2 trials, both run by Helsinn for gastrointestinal motility around 2008 to 2013, both completed, after which the program was discontinued [ClinicalTrials.gov]. The literature is tiny and shrinking, **54 lifetime publications and just 2 in 2024** [PubMed]. No company is pursuing it for the uses it is sold for.

Market. Market strength **82**, with roughly **504 listings across 203 sellers** [P08 market model; RUO price tracker]. The CJC-1295/ipamorelin blend is its own well-established product line.

Diffusion. Bottom-up and long-running. First social December 2018, one of the earliest research peptides to establish a social footprint [social listening].

Discourse. Mixed-to-positive, the most balanced of the growth-hormone peptides: of 34 reports, 10 positive, 17 neutral, 7 negative [social listening]. One user reported feeling *"great, highly motivated and libido was trending upward"* [social listening], while others flagged fatigue and throat sensations.

The read. Unproven but booming, in the growth-hormone wing. A pharmaceutically abandoned molecule with a thriving second life in the gray market, it exemplifies how the RUO economy resurrects compounds that industry has walked away from. Watch it as a stable core product rather than an emerging story.

38 · CJC-1295

What it is. A synthetic analog of growth-hormone-releasing hormone (GHRH), sold in two forms: "with DAC," a modification that extends its half-life to days, and "No DAC," a shorter-acting version favored for more physiological, pulsatile dosing. It is the standard GHRH partner to ipamorelin.

Evidence. Among the thinnest in the set. **A single registered trial** (a 2005 Phase 2 for HIV-related visceral obesity) that was **terminated**, and no other human trials [ClinicalTrials.gov]. PubMed holds just **33 lifetime publications and none in 2024** [PubMed]. The evidence base is close to empty.

Market. Correspondingly middling in the model, **59 (with DAC) and 57 (No DAC)**, the lowest market-strength scores of any research chemical profiled here, yet commercially deep: about **646 listings for the No-DAC form across 162 sellers**, plus 81 more for the DAC form and 69 for the ipamorelin blend [P08 market model; RUO price tracker].

Diffusion. Bottom-up. First social December 2018, alongside ipamorelin [social listening].

Discourse. Neutral-leaning, technical, and stack-oriented: of 24 reports, 9 positive, 13 neutral, 2 negative [social listening]. The talk is protocol design, one user explaining that *"No DAC feels more natural, pairs really well with Ipamorelin, and gives you way more control"* [social listening], the vocabulary of an experienced self-experimenter.

The read. Unproven but booming, with the widest gap in the set between near-zero evidence and deep commercial entrenchment. It is a product defined by its pairing rather than its own profile. Watch it as a bellwether for the growth-hormone-peptide category as a whole.

39 · Tesamorelin

What it is. An approved GHRH analog, sold as Egrifta for a narrow indication: reducing excess abdominal fat in people with HIV-associated lipodystrophy. Unlike its gray-market GHRH cousins, it has cleared regulatory review, which makes it the rare research-peptide-adjacent compound with genuine trial backing.

Evidence. A real program: **24 genuine trials spanning Phase 1 through Phase 4** (HIV, fatty liver disease, cognition), and **121 publications** [ClinicalTrials.gov; PubMed]. Approved, if for a niche use.

Market. Here the columns collide most productively. Tesamorelin scores **94 on market strength**, tied for third overall and level with tirzepatide, on about **532 listings across 205 sellers** [P08 market model; RUO price tracker], yet it draws only **81 news mentions** [news corpus]. An approved drug with blockbuster-tier commercial energy and almost no press.

Diffusion. Unusual: social slightly leads news, both appearing only recently (2024 to 2025) [social listening; news corpus], suggesting the wellness market discovered it as a "legitimate" GHRH option rather than the press covering it as medicine.

Discourse. Small but positive-leaning: of 11 reports, 4 positive, 5 neutral, 2 negative [social listening].

The read. A hybrid, straddling "proven but quiet" and "unproven but booming." The approval is real; the way it is being bought and used, off-label and unsupervised, for general body-composition and anti-aging goals, is not what the label authorizes. Watch it as the case where an approved niche drug gets absorbed into the wellness stack.

40 · PT-141 (bremelanotide)

What it is. A melanocortin-receptor agonist, approved by the FDA as Vyleesi for hypoactive sexual desire disorder in premenopausal women. Unlike drugs that act on blood flow, it works centrally, on brain pathways governing sexual desire. It is sold in the gray market to both sexes as a libido peptide.

Evidence. The most clinically validated non-GLP-1 peptide in the study. A **full Phase 1 to Phase 3 program, 10 genuine trials** run by Palatin, plus newer obesity and kidney studies, and **121 publications, 8 in 2024** [ClinicalTrials.gov; PubMed]. This is a real approved drug with real trials.

Market. Despite the strongest evidence in the research-peptide cohort, its market strength is a middling **73**, on about **373 listings across 180 sellers** with a search-intent signal near 120 [P08 market model; RUO price tracker]. Evidence and market energy are almost inversely related here.

Diffusion. Muted. Only 38 news mentions, and essentially no first-person social experience records tagged to it, despite steady commercial availability [news corpus; social listening]. Its use is real but private, discussed less openly than recovery or cognitive peptides.

The read. Proven but quiet, the mirror image of BPC-157. It holds the best evidence and the most modest momentum in the research-peptide world, a reminder that clinical validation and cultural traction are close to uncorrelated in this field. Watch it as the control case: what adoption looks like when the science is settled but the culture stays subdued.

41 · Epithalon

What it is. A synthetic four-amino-acid peptide (also spelled epitalon) from the Russian "Khavinson bioregulator" tradition, marketed for longevity and sleep on the claim that it activates telomerase, the enzyme that maintains chromosome ends. It is a research chemical with essentially no Western clinical presence.

Evidence. Zero registered trials on ClinicalTrials.gov, and **142 publications, only 1 in 2024**, most of it older Russian-language work that Western registries index poorly [ClinicalTrials.gov; PubMed]. The human evidence, such as it is, sits outside the systems this study can verify.

Market. Strong in the model despite the empty evidence column: market strength **89** and search intent near 200, on about **422 listings across 180 sellers** [P08 market model; RUO price tracker].

Diffusion. Bottom-up and social-native. First social July 2020, and a news presence so slight the diffusion tracker records it as effectively zero [social listening; news corpus].

Discourse. The most uniformly positive profile in the set: of 21 reports, 8 positive, 13 neutral, and **zero negatives** [social listening]. One user described waking *"feeling like I had gotten quality rest"* even on fewer hours [social listening]. A total absence of negative reports, for a compound with no controlled evidence, is the frictionless-testimony pattern in its purest form.

The read. Unproven but booming, the anchor of the longevity-bioregulator wing. It runs almost entirely on belief and a Russian research tradition that Western science has neither replicated nor refuted. Watch it as the compound whose adoption is least connected to any verifiable evidence at all.

42 · Semax

What it is. A Russian heptapeptide derived from a fragment of ACTH, approved in Russia as a nootropic and stroke therapy but unapproved in the West. It is marketed for focus, cognition, and "brain fog," and anchors the overlap between the peptide and nootropics subcultures.

Evidence. Zero registered trials on ClinicalTrials.gov, its clinical work living in Russian registries, and **231 publications, 4 in 2024** [ClinicalTrials.gov; PubMed]. Like epithalon, its evidence base is real-but-unverifiable from Western sources, which is a different situation than "no evidence" but leaves the same gap for a US buyer.

Market. Market strength **81**, search intent near 180, on about **464 listings across 189 sellers** at a low median near $5.36 per milligram [P08 market model; RUO price tracker].

Diffusion. Bottom-up. First social July 2019, first Western news not until 2025 [social listening; news corpus].

Discourse. Cognition-focused and genuinely mixed: of 31 reports, 7 positive, 16 neutral, 7 negative [social listening], one of the more balanced research-peptide profiles. A representative account described *"clarity, the mental static dropping, easier task initiation, less procrastination"* [social listening], the vocabulary of the productivity-optimization scene.

The read. Unproven but booming, at the frontier where the peptide movement meets nootropics, a subculture with its own logic and its own tolerance for foreign-approved, Western-unapproved compounds. Watch it as the bridge between two adjacent gray markets.

43 · 5-Amino-1MQ

What it is. The outlier of the dossier: not a peptide at all, but a small molecule that inhibits an enzyme called NNMT, pursued for fat loss and metabolic health. It is included because the wellness market sells, stacks, and discusses it exactly as if it were one of the peptides.

Evidence. The thinnest in the entire study. **Zero human trials and just 3 lifetime publications**, entirely preclinical [ClinicalTrials.gov; PubMed]. There is essentially no human data of any kind.

Market. Absent from the market-strength model, yet commercially real: about **365 listings across 151 sellers** at a low median near $4.40 per milligram [P08 market model; RUO price tracker]. It is sold at scale despite having almost no scientific existence.

Diffusion. Social-native. First social July 2020, with only a trace news presence (10 mentions) [social listening; news corpus]. No first-person experience records were cleanly extractable, consistent with a compound whose culture is still thin.

The read. Unproven but booming, and the purest illustration of the report's central inversion: a commodity market with dozens of sellers has assembled around a molecule that three papers and no trials describe. Watch it as the limiting case, the point at which commercial availability has decoupled from scientific validation almost completely.

"Watch" throughout this dossier marks where the diffusion dynamics are most active or most instructive. Nothing in Part VII is a recommendation to purchase or use any compound. Many are unapproved research chemicals of unverified purity and identity;

their inclusion is analysis of a phenomenon, not an endorsement of it.

PART VIII

Trajectories

44 · Where each world goes next

Forecasting is inference, not measurement, and this part is labelled accordingly. Nothing that follows is a finding in the data. It is the report's reasoned reading of where the observed trends point, offered with the same discipline that separates "people are reporting X" from "X has been shown." Where the evidence supports a confident expectation, the report says so; where the outcome turns on variables outside this dataset, it names them rather than guessing.

The single most useful frame is the quadrant map from Part VII. The peptide phenomenon is two stories moving at different speeds: a proven-and-popular pharmaceutical franchise, and a booming-but-unproven consumer market. They face different futures, and the interesting questions are different for each.

45 · The GLP-1 institutionalization path

Inference, well-supported. The proven-and-popular quadrant is expanding, not cresting. Three forces point the same direction. The pipeline is deepening: retatrutide's triple-agonist trials are producing the largest weight reductions yet reported [ClinicalTrials.gov], and its position atop the commercial model (98/100) shows demand already assembled ahead of approval [P08 market model]. The delivery problem is being solved: orforglipron's 25 Phase 3 trials point toward an oral option that removes the injection barrier and, with it, the largest source of friction and discontinuation [ClinicalTrials.gov]. And the indications are broadening: the publication curve, more than 2,500 GLP-1 papers in 2024 alone [PubMed], increasingly covers cardiovascular, liver, and early addiction and cognition questions rather than diabetes and obesity alone.

The binding constraint on this path is supply, price, and access, not demand. That is the profile of a durable franchise settling in for a long run rather than a bubble. For an investor, the strategic questions are no longer "is this real" but "how does the market split between injectable and oral, single and multi-hormone, branded and eventually generic," and "which secondary indications convert from hypothesis to label." The measurement caution from Part Four applies with full force: the "Ozempic changes everything" spillover narratives, drinking, groceries, air travel, remain claims with zero established causal backing in the corpus [news corpus; social listening], and a disciplined reader should treat the societal story as culture, not yet as demonstrated economics.

46 · The research-peptide reckoning

Inference, genuinely uncertain. Everything in this report says the non-GLP-1 movement is real, growing, and commercially entrenched, dozens of sellers, hundreds of listings per compound, an established subculture with its own vocabulary and norms. But it lives almost entirely in the

unproven-but-booming quadrant, and that is not a stable place to stay. A market this large standing on a clinical base this thin will be resolved, one way or another. Three scenarios are visible. They are not mutually exclusive, and different compounds may travel different roads.

Scenario one: legitimization. A handful of compounds with the strongest anecdotal and preclinical cases attract genuine trials and either cross toward approved use or are ruled out. The leading edge already exists: nearly every *first real* efficacy trial of BPC-157, GHK-Cu, MOTS-c, and melanotan II is a single Phase 2 registered by one sponsor, Hudson Biotech, all beginning in early 2026 [ClinicalTrials.gov]. BPC-157 is the obvious bellwether. If its 2026 Phase 2 reads out positive, it becomes the proof-of-concept that online popularity can pull a compound into legitimate development; if it fails, it becomes the cautionary tale. Either outcome sorts the field, because for the first time there will be human evidence where there is now only anecdote.

Scenario two: regulatory contraction. Growing regulatory attention pushes the gray market underground. The FDA has already acted against several of these compounds in the compounding context, and the price tracker's own signals, two sellers flagged for suspected review manipulation, sixty of eighty-six sellers making no manufacturing-origin claim at all [RUO price tracker], describe exactly the weak-trust, opaque-provenance environment that invites a crackdown. In this scenario the visible seller base thins, activity shifts to harder-to-observe channels, and the measurable market shrinks even if underlying use does not.

Scenario three: continued gray-zone growth. Absent a decisive safety event or enforcement wave, the market simply keeps expanding in its current legally ambiguous form, with commerce and culture continuing to outrun both science and regulation. This is the path of least resistance and, on current evidence, the modal outcome: the infrastructure is built, the demand is demonstrated, and nothing in the data suggests an imminent forcing event.

Which scenario dominates depends on variables outside this dataset: whether a high-profile safety incident occurs, how aggressively regulators choose to act, and whether any research peptide accumulates enough real evidence to change its status. The honest position is that all three are live, and the next eighteen months of trial readouts and regulatory posture will tilt the balance.

47 · The measurement lesson

Inference, strongly supported. The most durable finding here is not about any compound but about *where attention lives*. For an entire class of consumer-health products, the mainstream record, the news, lagged the actual phenomenon by years. BPC-157 was a fixture of social media from late 2019 and did not enter the news corpus until mid-2024 [social listening; news corpus]; GHK-Cu, MOTS-c, epithalon, and 5-Amino-1MQ show the same shape. The reason is structural: these movements did not generate the events, an FDA approval, an earnings surprise, a celebrity confession, that newsrooms are built to detect. They generated something quieter, millions of TikTok views and a lengthening list of vendors, which the institutional record does not register until much later.

The implication generalizes well beyond peptides. Anyone trying to understand emerging consumer-health behavior by reading coverage of it will be systematically late, because coverage is a lagging indicator of exactly this kind of phenomenon. The culture shows up first in social platforms and commercial listings; the institutions arrive afterward. For an analyst, the actionable version is simple: to see the next wave early, watch the leading indicators, not the newspaper.

48 · Leading indicators worth tracking

If the news lags, what leads? The evidence base points to a short list of signals that moved *before* mainstream recognition, and that would move first again:

- **Social first-appearance and sustained posting**, especially on Reddit, where the technical, protocol-level conversation begins. A compound accumulating dosing threads and "week three update" posts is being used well before it is covered.
- **Seller and listing proliferation** in the RUO price corpus. New sellers do not stock what no one buys; a compound climbing from a handful of listings toward a hundred is a demand signal that precedes any headline.
- **The market-strength and search-intent scores** as early commercial tells. Retatrutide topping the model at 98 while still unapproved [P08 market model] is precisely the kind of pre-recognition signal that pays attention.
- **First company-run trials**, such as the Hudson Biotech 2026 cluster [ClinicalTrials.gov], which mark the moment a gray-market compound begins its possible migration toward legitimacy.

Tracked together, these turn the lag into an advantage: they are visible now for compounds the press will not cover for years.

49 · What would sharpen this picture

The report has been careful to flag what its evidence cannot do, and three additions would materially reduce that uncertainty. Each targets a specific blind spot.

- **An official search-interest series.** Direct Google Trends data would convert the search-intent proxy into a real, time-resolved measure of public curiosity, allowing month-to-month trend claims the report currently declines to make and cleanly separating "curiosity" from "intent to buy." This is the single highest-value addition a future edition could make.
- **Real-world usage or utilization data.** Prescription, claims, or survey data would convert "there is a market" into "this many people," the one question the current evidence base genuinely cannot answer for the research-peptide world, where no such instrument exists.
- **The first genuine efficacy trials reading out.** The 2026 Hudson Biotech Phase 2 studies of BPC-157, GHK-Cu, MOTS-c, and melanotan II [ClinicalTrials.gov] will, for the first time, place controlled human evidence beside the anecdote for specific compounds, collapsing the science-versus-culture gap where it matters most.

Until those exist, the responsible summary is the one this report has held throughout. The peptide moment is real, it is plural, and most of it, outside the GLP-1 franchise, is still a story of what people are doing and hoping rather than of what has been shown. Keeping those two sentences apart is the entire discipline of this study, and it is also the most reliable guide to what happens next.

Apparatus

50 · Methodology

What this study is. A descriptive synthesis of a large, pre-existing evidence base assembled by the therapeutic-peptide adoption research project. It compiles, analyzes, and interprets collected data; it does not run new experiments, and it makes no clinical claims. The unit of analysis is the *compound* and the *signal* (news, social, commercial, and so on), not the patient.

Primary data sources.

- **News corpus**, roughly 32,700 articles (about 26,900 with full text acquired) in a PostgreSQL journalism database, discovered and enriched through a mix of licensed news APIs, RSS, and web extraction, spanning April 2015 to August 2026. Each article is scored for quality, de-duplicated, tagged by topic, and linked to the compounds it mentions.
- **Social corpus**, roughly 18,100 compound-linked posts across TikTok, Instagram, and Reddit, spanning 2013 to August 2026, with engagement metrics and, where extractable, structured "experience" records.
- **Commercial price corpus**, a live tracker of U.S. research-use-only peptide retailers (1,000+ candidate domains, 443 verified, 302 with captured live prices, ~18,400 listings, 114 compounds), collected under a strict robots-and-terms-compliant, no-evasion policy.
- **Market-strength model**, a 0–100 per-compound score blending commercial shelf presence and search-intent signals, imported from the companion market project. Explicitly *not* audited pharmaceutical revenue.
- **Public registries**, ClinicalTrials.gov and PubMed, queried directly to characterize clinical development and scientific-publication activity (Sections 06–07).

How the compound taxonomy works. Compounds and their aliases come from a canonical catalog of 301 tracked molecules, each carrying a functional category and a GLP-1 flag. Mentions are matched to compounds by canonical name and alias, which drives every share-of-voice and diffusion measure.

The GLP-1 control. Because two GLP-1 drugs dominate every raw signal, every major measure is computed three ways where possible, all compounds, GLP-1 only, and non-GLP-1 only, so that the broader field can be seen with the dominant class held constant. This control is the analytical backbone of the report's central finding.

Share of voice and diffusion. For cross-compound comparison the report uses *share of voice* (a compound's mention count within a medium) rather than raw totals, because it is more robust to how much of each medium was collected. *Diffusion signatures* are built from each compound's first-appearance and peak months in news versus social, which is how the top-down and bottom-up pathways are distinguished.

A caveat that shapes every temporal claim. The news corpus was *harvested* in 2026, and news APIs return recent material far more readily than old material. Raw monthly article counts are therefore a artifact of *when the corpus was built* as much as of when attention actually occurred, and the report never treats a raw monthly spike as proof of a real attention spike. Trend claims rest

instead on ratios (GLP-1 versus non-GLP-1 within the same month, which share the same collection bias), on share of voice, and on social post timing, which reflects real posting dates back to 2013. This is the most important methodological guardrail in the study.

51 · Limitations

This report is only as good as its evidence, and the evidence has real limits. Stated plainly:

- **No usage or prescription data.** The study cannot count users of anything. It observes markets, media, and discourse, not consumption. For research peptides specifically, no reliable population-use estimate exists anywhere, and this report does not manufacture one.
- **No official search series.** Google Trends data was not available; search interest is approximated by commercial-intent and social proxies, which cannot show clean month-to-month curiosity trends (Section 10).
- **Harvest recency bias.** Raw news volume over time is distorted by when the corpus was collected; only ratio- and share-based temporal claims are treated as reliable (Section 19).
- **Social data is a non-representative convenience sample.** The social corpus captures public posts that were discoverable and matchable, not a representative sample of any population. Engagement metrics are platform-specific and not directly comparable. Absence of a compound from the corpus is not proof of absence from the culture.
- **Experience and sentiment data are machine-extracted anecdotes.** The 1,917 experience records are LLM-parsed self-reports, carry a standing "not a clinical outcome" disclaimer, and are subject to extraction error, selection, and self-report bias. They evidence discourse, never efficacy.
- **Cultural-spillover tags are claims, not measurements.** The 8,400+ GLP-1 effect observations record what media and social sources *assert*, not verified transactional or clinical effects; the report treats them strictly as discourse.
- **Syndication and press-release detection are heuristic.** De-duplication and press-release flags are automated approximations that would benefit from human validation; they are used to *discount* raw counts, which is the conservative direction.
- **Commercial coverage is a floor, not a census.** The price corpus deliberately excludes bot-walled and login-gated sellers and covers a short collection window, so it undercounts the true market and cannot yet report price *trends* reliably.
- **Advertising is largely unmeasured.** The ad stream is characterized qualitatively; a creative-level ad audit was not performed for this edition (Section 13).
- **Geography is weak.** The data cannot support a reliable geographic or demographic breakdown of interest; claims about *who* and *where* are intentionally limited.
- **This is a snapshot.** Every number reflects the corpus as of August 2026 and will drift as the underlying data grows.

None of these limitations undercut the report's central, structural findings, the news/social split, the two diffusion pathways, and the science-versus-culture gap are large and robust to all of them. They do mean that precise magnitudes, and especially trends over time, should be read as directional rather than exact.

52 · How evidence was labelled

The report leans on an explicit hierarchy of evidence, and its prose is written to keep the reader aware of which kind supports any given claim. From strongest to weakest for the purpose of establishing that something is *true of the world*:

1. **Administrative / utilization data** (claims, prescriptions, official statistics), authoritative for real-world use. *Largely unavailable to this study; its absence is disclosed, not papered over.*
2. **Scientific evidence** (peer-reviewed literature, clinical trials, systematic reviews), establishes what a compound does. Sourced here from PubMed and ClinicalTrials.gov.
3. **Commercial evidence** (listings, prices, seller counts, market models), establishes that a market exists. Strong in this study.
4. **Search behavior**, an early, honest signal of curiosity. Proxied, not measured directly.
5. **Social behavior** (posts, engagement, community norms), establishes discourse and culture. Rich in this study, but evidence of *conversation*, not *fact*.
6. **Journalism**, establishes what has entered mainstream conversation, discounted for syndication.
7. **Anecdotal evidence** (individual experience reports), establishes what people *say* happened. The weakest class for establishing truth, the strongest for understanding motivation.

The underlying database enforces a parallel discipline at the row level, tagging every derived observation by data class, raw observation, deterministic derivation, model/AI classification, human-curated, or statistical estimate, so that machine-generated labels are never silently promoted to facts. The single rule that governs the entire report follows from this hierarchy: **"people are reporting X" and "X has been demonstrated" are different sentences, and the report never lets one stand in for the other.**

53 · Glossary, key terms in plain language

A quick reference for the non-specialist. Terms are defined as this report uses them.

- **Peptide.** A short chain of amino acids, the same building blocks that make up proteins, just fewer of them linked together. The boundary with "protein" is a matter of size and convention, not a hard line.
- **Amino acid.** One of the ~20 small molecules that link in sequence to form peptides and proteins. The order of amino acids determines what the molecule does.
- **GLP-1 (glucagon-like peptide-1).** A hormone the gut releases after eating. It prompts insulin release, slows stomach emptying, and signals fullness to the brain. The blockbuster weight-loss drugs are engineered, long-lasting copies of it.
- **GIP, glucagon.** Two related metabolic hormones. Tirzepatide targets GIP as well as GLP-1 (a "dual agonist"); retatrutide adds glucagon (a "triple agonist"). Hitting more targets appears to deepen the effect.
- **Incretin.** The class of gut hormones (including GLP-1 and GIP) that link eating to insulin release. "Incretin drugs" is a precise umbrella term for the GLP-1 family.
- **Agonist.** A molecule that switches a receptor on, mimicking the body's own signal. (An antagonist blocks it.)

- **Half-life.** How long a drug lasts in the body before half of it is cleared. Engineering a peptide to resist breakdown is what turned a minutes-long natural hormone into a once-weekly injection.
- **Peptide bioregulator.** A short peptide, often from the "Khavinson" Russian research tradition (for example epithalon), claimed to regulate the function of a specific tissue. Human evidence is generally thin.
- **RUO, "research use only."** A label allowing a chemical to be sold for laboratory research, not for human use. Much of the online peptide market operates under this status: legal to sell as a reagent, not approved as a medicine, but bought and self-administered anyway.
- **Compounding pharmacy.** A licensed pharmacy that mixes a drug to order, for example to fill a shortage or personalize a dose. A regulated channel that expanded sharply during the GLP-1 shortage.
- **Telehealth.** Remote medical care, here the online clinics that prescribe and ship GLP-1 and other peptide therapies. The visible, regulated cousin of the gray market.
- **Off-label.** Prescribing an approved drug for a use its label does not cover. Legal for physicians; distinct from the unapproved research-chemical market.
- **Clinical trial phases.** Phase 1 tests safety in a few people; Phase 2 tests whether it works and at what dose; Phase 3 is the large trial that supports approval. A registered trial is development *activity*, not proof of anything yet.
- **ClinicalTrials.gov / PubMed.** The public U.S. registry of human trials, and the public index of biomedical literature. This report uses trial counts as a measure of development activity and publication counts as a measure of research activity, neither of which is a measure of quality or of real-world use.
- **Share of voice.** A compound's share of all mentions within one medium (news or social). Used instead of raw counts because it is more robust to how much of each medium was collected.
- **Diffusion pathway.** The route a compound travels into public life. "Top-down" runs science to trials to press to public; "bottom-up" runs community to social media to online vendors to belated coverage.
- **Syndication.** The republishing of one article across many outlets. Inflates raw coverage counts, so this report discounts for it; syndication is not independent attention.
- **Market-strength model (P08).** An internal 0 to 100 score blending a compound's commercial shelf presence and search-intent signals. A proxy for commercial momentum, explicitly *not* audited pharmaceutical revenue.
- **Value chain.** The layered set of businesses that turn a molecule into a sold product: innovators, manufacturers, compounders, telehealth platforms, vendors, and testing labs. Where in this chain the durable margin sits is a central investor question.
- **Certificate of analysis (COA).** A lab document attesting a product's identity and purity. In an unregulated market it is the real quality gate; a low price without a COA is not a bargain.
- **Evidence class.** The kind of evidence behind a claim: administrative, scientific, commercial, search, social, journalistic, or anecdotal. This report labels the class of every material claim and never lets a weaker class stand in for a stronger one.

54 · Sources and provenance

This report is a synthesis of collected datasets and public registries rather than a literature review, so its "sources" are primarily data corpora. Each is named with its role and its known limits. Specific scientific and clinical claims in Sections 06–07 also cite the underlying registry records inline.

Corpus and model sources

1. **Therapeutic-peptide journalism corpus**, PostgreSQL `peptide_news_07` (RP-50 / SP-07). ~32,700 articles, ~26,900 with acquired full text, April 2015 – August 2026; topic tags, de-duplication, compound linkage, press-release heuristics. *Journalism and its analytical layer; syndication is heuristic.*
2. **Social-listening corpus**, TikTok, Instagram, Reddit; ~18,100 compound-linked posts, 2013 – August 2026, with engagement and extracted experience records. *Non-representative public-post sample; discourse evidence only.*
3. **GLP-1 effect-observation layer**, ~8,400 tagged media/social claims of GLP-1 secondary and tertiary effects. *Claims, not verified effects; no established causal evidence in any category.*
4. **RUO commercial price tracker**, SQLite retailer/price corpus, ~18,400 listings across 302 priced U.S. sellers, July–August 2026, collected under a robots-and-terms-compliant, no-evasion policy. *Commercial availability; a floor on true market size.*
5. **Market-strength model (P08)**, per-compound 0–100 score blending commercial shelf and search-intent signals. *A proxy; explicitly not audited pharmaceutical revenue.*
6. **Canonical compound catalog**, 301 tracked compounds with functional categories, aliases, and GLP-1 flags, sourced from the therapeutic-peptide research library.
7. **ClinicalTrials.gov** (API v2), trial registrations by compound, queried for Section 07. *Registrations are development activity, not approval or use.*
8. **PubMed / NCBI**, publication counts by compound and year, queried for Section 06. *Publication activity, not quality or human relevance.*

Registered but not bulk-analyzed in this edition (available for future extension): a collected peptide social-advertising manifest; a therapeutic-peptide YouTube library; a peptide-ecosystem observatory dataset; and a large-scale GLP-1 secondary-effects monitor. These are named for transparency and provenance, not relied upon for specific claims here.

Public press. Named outlets (CNBC, Yahoo Finance, the *Guardian*, the *Daily Mail*, Medscape, *Nature*, and others) are characterized in aggregate as the composition of the news corpus; no third-party article text is reproduced.

Coverage and provenance statement. The findings in this report are bounded by the corpora above as they stood in August 2026. The report is designed to be shareable: it contains market, media, and cultural analysis and no confidential commercial, cost, or margin data. Where a data stream was unavailable (utilization data, official search series) or unmeasured (advertising creative, geography), the report states the gap rather than substituting a weaker signal presented as a stronger one. All quantitative figures are reproducible from the analysis scripts and the evidence JSON archived with this report.

55 · Appendix, reference tables

A. Most-covered compounds in the news (share-of-voice mention count; GLP-1 flagged)

RANK	COMPOUND	GLP-1?	NEWS MENTIONS
1	Semaglutide	GLP-1	17,186
2	Tirzepatide	GLP-1	12,291
3	Insulin	—	3,265
4	Orforglipron	GLP-1	2,786
5	Liraglutide	GLP-1	2,702
6	Retatrutide	GLP-1	2,153
7	Dulaglutide	GLP-1	1,469
8	CagriSema	GLP-1	1,324
9	BPC-157	—	1,098
10	Metformin	—	600

B. Most-discussed compounds on social media (share-of-voice mention count)

RANK	COMPOUND	REGULATORY REALITY	SOCIAL MENTIONS
1	GLP-1 (as a class term)	—	505
2	Insulin	Approved	345
3	GHK-Cu	Cosmetic / research	299
4	Semaglutide	Approved	286
5	BPC-157	Research chemical	274
6	NAD+	Supplement / research	267
7	GHK	Cosmetic / research	261
8	Tirzepatide	Approved	244
9	TB-500	Research chemical	206
10	IGF-1	Research / restricted	189

C. Diffusion signatures (first appearance and dominant medium; "lag" = social leads news)

COMPOUND	FIRST SOCIAL	FIRST NEWS	PATHWAY
Semaglutide	2019	2019	Top-down (news-led)
Tirzepatide	2022	2021	Top-down (news-led)
Retatrutide	2024	2022	Top-down, pre-approval
BPC-157	2019	2024	**Bottom-up (5-yr lag)**
TB-500	2021	2024	Bottom-up
MOTS-c	2020	2025	Bottom-up
GHK-Cu	2020	2025	Bottom-up (cosmetic origin)
KPV	2020	—	Social-only

D. Commercial market-strength leaders (0–100 model; note the research chemicals near the top)

COMPOUND	SCORE	EVIDENCE STATUS
Retatrutide	98	Investigational
Tirzepatide	94	Approved
Tesamorelin	94	Approved (niche)
MOTS-c	93	Research chemical
BPC-157	93	Research chemical
GHK-Cu	93	Research chemical
Semaglutide	90	Approved
Epithalon	89	Research chemical

STANDING CONSTRAINT

This document is a study of markets, media, and culture. It describes what is sold, searched, discussed, and reported about therapeutic peptides. It does **not** recommend the use of any compound in any person, specifies no dose, route, or schedule, and is not medical advice. Many compounds discussed here are unapproved research chemicals of unverified purity and identity; their appearance in this report is analysis of a phenomenon, never an endorsement of it.

SOUTH BEACH LONGEVITY · Peptide Market & Culture Intelligence · Science, Optimization, Longevity. *From the Lab to the Feed* · compiled August 2026. A market, media, and cultural analysis of collected news, social, commercial, and public-registry data. It describes what is researched, sold, searched, discussed, and reported about therapeutic peptides. It recommends no human use of any compound, specifies no dose, route, or schedule, and is not medical advice.