



The Age of Therapeutic Abundance

PHARMACEUTICAL SCIENCE IS ADVANCING
FASTER THAN THE AMERICAN HEALTHCARE SYSTEM



813

813 Strategies

2026

Executive Summary

An investigational obesity drug has now produced average weight loss approaching 30% in Phase III. CRISPR-edited stem cells are approved medicine. A cancer treatment manufactured around the mutational fingerprint of an individual patient has produced a positive Phase III result. Researchers are testing whether autoimmune diseases can be pushed into durable drug-free remission, whether muscle can be preserved while fat disappears, and whether artificial intelligence can find useful drug targets and molecules faster than conventional discovery.

These advances come from different technologies and operate against different diseases, but they are arriving at roughly the same time. Medicine is gaining more ways to intervene upstream in disease, target the biology of an individual patient, or alter the biological process itself.

Obesity medicine makes the scale easiest to see. Semaglutide produced 14.9% mean weight loss after 68 weeks in the 1,961-patient STEP 1 trial. Tirzepatide moved mean reductions beyond 20% in SURMOUNT-1. Lilly reported in May 2026 that participants receiving 12 milligrams of the investigational drug retatrutide in Phase III lost an average of 70.3 pounds, or 28.3% of body weight, over 80 weeks; 45.3% lost at least 30% of their starting weight. Among participants with baseline BMI of at least 35 who continued into a 104-week extension, average loss reached 30.3%.

Obesity sits upstream from an enormous amount of chronic disease. Measured national data show obesity in 40.3% of American adults and overweight in another 32.1%. CDC estimates that 40.1 million Americans had diabetes in 2023 and 115.2 million adults had prediabetes. SELECT has already shown that semaglutide can reduce the combined risk of cardiovascular death, nonfatal myocardial infarction or nonfatal stroke by 20% in adults with overweight or obesity and established cardiovascular disease but without diabetes.

Researchers are now trying to manipulate what kind of tissue disappears during aggressive weight loss. In a randomized Phase II trial published in *Nature Medicine* in June 2026, the anti-myostatin antibody apitegromab preserved 54.9% of the lean mass that otherwise would have been lost during tirzepatide treatment. Regeneron is pursuing a similar strategy with semaglutide and its own muscle-preservation combinations. Weight loss is beginning to become a problem of body composition rather than a single number on a scale.

Scale is developing alongside efficacy. FDA approved orforglipron, marketed as Foundayo, in April 2026 as a once-daily oral small-molecule GLP-1 receptor agonist for chronic weight management. Its pivotal Phase III trial produced 11.1% mean weight reduction at 72 weeks at the highest studied dose. That is less than the best injectable therapies, but a pill that is easier to manufacture, distribute and take can serve a very different market. Tens of millions of patients are unlikely to be treated with one drug at one dose through one delivery system.

Cancer medicine is moving in the opposite direction, toward increasingly individualized treatment. On August 19, 2026, Merck and Moderna announced that intismeran autogene, an mRNA therapy manufactured around the mutational fingerprint of each patient's tumor, met both its recurrence-free-survival and distant-metastasis-free-survival endpoints in a randomized Phase III melanoma trial involving 1,137 patients. Each product can encode as many as 34 neoantigens selected from that patient's cancer.

Nine days later, BioNTech terminated a Phase II trial of another individualized mRNA treatment in resected colorectal cancer after an independent monitoring board concluded that continued treatment was unlikely to change the efficacy outcome and identified a numerical imbalance in overall survival.

Those two results belong together. Personalized cancer medicine has advanced far enough that different products, combinations and tumor types are succeeding and failing in meaningful human trials. Patient-specific pharmaceutical manufacturing has moved beyond a laboratory demonstration.

Cell and gene therapy have crossed an even more basic threshold. Casgevy became the first FDA-approved CRISPR-based treatment in 2023, and in July 2026 its indication was expanded to patients age two and older with sickle cell disease with recurrent vaso-occlusive crises or transfusion-dependent beta thalassemia. The treatment edits a patient's own blood-forming stem cells to increase fetal hemoglobin production.

The science is extraordinary. The delivery is difficult. Patients require cell collection, specialized manufacturing, myeloablative conditioning, hospitalization and long-term follow-up. Casgevy entered the market at roughly \$2.2 million per patient, while Lyfgenia, a different gene therapy for sickle cell disease, was listed at roughly \$3.1 million.

The financing problem was foreseeable enough that CMS had already selected a Medicaid-focused Cell and Gene Therapy Access Model for testing in February 2023, before either therapy reached the U.S. market. The model now negotiates outcomes-based arrangements on behalf of participating Medicaid programs. Thirty-two states, the District of Columbia and Puerto Rico are participating, representing approximately 84% of Medicaid beneficiaries with sickle cell disease.

The medicine is forcing insurance accounting to operate on a different clock.

The same pressure is beginning to appear elsewhere. A 2024 New England Journal of Medicine case series followed 15 patients with severe refractory autoimmune disease after CD19 CAR-T treatment. All eight patients with lupus met the study's remission criterion, while patients with inflammatory myositis and systemic sclerosis achieved major clinical improvements; immunosuppressive therapy was discontinued. Fifteen highly selected patients cannot establish a standard of care, but they are enough to ask a different question: whether selected immune diseases can sometimes be reset deeply enough to produce durable remission rather than managed indefinitely with recurring immunosuppression.

Artificial intelligence belongs in this story more narrowly than the hype sometimes suggests. FDA's Center for Drug Evaluation and Research had experience with more than 500 submissions containing AI components from 2016 through 2023. Rentosertib, whose TNIK target and small-molecule design were developed through an AI-directed platform, produced a preliminary signal in a 71-patient randomized Phase IIa study in idiopathic pulmonary fibrosis and has since advanced into Phase III; the first patient in the GENESIS-IPF-3 trial was dosed on September 10, 2026. Computational target discovery and molecule design have moved from software demonstrations into late-stage human drug development while remaining subject to ordinary pharmaceutical uncertainty.

The scientific advance is arriving inside an access system already struggling with ordinary medicine. A KFF analysis of CMS reporting data found that Medicare Advantage insurers made approximately 52.8 million prior-authorization determinations in 2024 and denied roughly 4.1 million requests in whole or in part. Only 11.5% of denials were appealed, but 80.7% of appealed denials were overturned. The 2025 AMA Prior Authorization Physician Survey found an average of 40 prior authorizations per physician each week, consuming roughly 13 hours of physician and staff time.

CMS is trying to modernize that process. Its 2024 rule imposed new electronic and timing standards for many medical services; in April 2026, the agency proposed separate pharmaceutical requirements that would impose electronic standards and, in several federal programs, maximum decision periods of 24 hours for urgent requests and 72 hours for standard requests beginning in 2027.

American laboratories can sequence a tumor and manufacture an individualized therapy around its mutations while the country is still building uniform electronic rules for how an insurer approves the prescription.

The financing problem runs deeper. Beginning July 1, 2026, Medicare's GLP-1 Bridge began providing qualifying Part D beneficiaries certain obesity medicines for a \$50 monthly copayment through a national demonstration operating outside ordinary Part D payment. Private employers are moving in the other direction in some cases. Forty-nine percent of large employers covered GLP-1 drugs for weight loss in 2025, but 6% dropped coverage in 2026, another 5% planned or were considering doing so for 2027, and 27% tightened or planned to tighten utilization controls.

A therapy can work and still be difficult to finance. The payer buys the drug today; the avoided heart attack, dialysis episode, disability claim or cancer recurrence may come years later, perhaps after the patient has moved to another insurer.

For much of modern medicine, finding something that worked was the dominant constraint. That constraint is beginning to loosen in several important fields. The pressure is moving downstream into manufacturing, diagnostics, financing, insurance rules, treatment capacity and the time required to get an eligible patient through all of them.

The United States does not need to approve every drug, pay every price or eliminate scrutiny. Failed drugs should fail. Weak evidence should remain weak evidence. Clinically interchangeable products should compete aggressively on price.

Once strong evidence exists and a patient clearly qualifies, however, the system has to be capable of acting on that information.

Therapeutic abundance will matter only if medicine can actually reach the people it was invented for.

I. The Evidence for Therapeutic Abundance

Biomedical history supplies plenty of reasons to distrust declarations of a golden age. Small studies fail in larger populations, surrogate endpoints exaggerate practical benefit, attractive mechanisms collapse in human trials, and biotechnology companies operate in capital markets that reward promising early results.

The useful dividing line is human evidence. FDA approval and completed randomized Phase III trials in clinically relevant populations carry the most weight. Randomized Phase II evidence can establish meaningful biological or clinical signals, while early human studies usually establish feasibility more than comparative effectiveness. Preclinical work, uncontrolled anecdotes and gray-market consumer use can point toward a direction without deserving the same confidence.

Several fundamentally different technologies have now crossed meaningful clinical thresholds at roughly the same time. Genome editing is FDA-approved medicine. Autologous cellular treatments have been approved in hematologic and solid tumors. Engineered T-cell receptors have entered clinical use. Radioligands, antibody-drug conjugates and bispecific antibodies have become recurring therapeutic categories. Modern obesity drugs have moved from weight endpoints into cardiovascular outcomes, and individualized mRNA cancer treatment has produced a positive Phase III readout. AI-directed molecular discovery remains much earlier, although it has now reached Phase III testing.

Many of these technologies are platforms. CRISPR can edit different genetic targets. CAR-T can redirect immune cells toward different antigens. mRNA can encode different biological information. Antibodies can block a receptor, carry a payload or connect two cells. Radioligands combine a targeting system with radioactive material, while computational discovery can be applied repeatedly across different biological targets.

A failed drug built on one of these systems does not necessarily return the entire technology to scientific irrelevance. Failure can reveal where a platform works, where it does not and what needs to change next.

Hospitals and insurers can handle an occasional unusual treatment through an exception, a special reimbursement code or a small network of specialty centers. Repeated exceptions eventually stop being exceptions. They become infrastructure.

That transition has begun.

II. Oncology and the Transition From Disease Categories to Individual Biology

Cancer treatment has been moving toward biological specificity for decades. Anatomical location still matters, but many treatment decisions now turn on whether a tumor expresses a receptor, carries a mutation, presents an antigen or has an immune environment that makes a particular therapy more or less likely to work.

Once treatment eligibility depends on a genomic assay, protein-expression test or molecular scan, the diagnostic is part of the treatment pathway even if an insurance system processes the test and the drug as separate claims.

Cellular Therapy Moves Into Solid Tumors

CAR-T showed that a patient's immune cells can be collected, genetically modified into a therapeutic product and returned to the body with profound effects in some blood cancers. The process remains demanding: T cells are collected, modified to express a receptor that recognizes a cancer target, expanded, reinfused after lymphodepleting therapy and monitored for potentially severe immune and neurologic complications.

Solid tumors have been harder because their targets are less uniform and their surrounding tissue can suppress the immune response. Cellular treatment has nevertheless begun to cross that divide.

FDA approved lifileucel in February 2024 for adults with unresectable or metastatic melanoma after prior systemic therapy, the first approved tumor-infiltrating lymphocyte treatment. Lymphocytes are taken from the patient's tumor, expanded outside the body and reinfused after lymphodepletion. The pivotal evidence came from a single-arm study rather than a randomized comparison, yet the approval established autologous cellular therapy as a regulated treatment for a solid malignancy.

Tecelra added another proof point. FDA approved the engineered T-cell receptor therapy for eligible patients with unresectable or metastatic synovial sarcoma whose tumors express MAGE-A4 and who carry certain HLA types.

The disease label is only the first screen. Treatment depends on the patient's tumor antigen, HLA genotype and clinical fitness to complete cell collection, manufacturing and therapy. Similar fragmentation already appears throughout targeted oncology and will expand as sequencing, antigen expression and individualized manufacturing determine who can receive a drug.

Individualized Neoantigen Therapy

The Merck-Moderna intismeran program pushes personalization further because the pharmaceutical product itself changes from patient to patient. Tumor material is sequenced, mutations are identified, computational systems select candidate neoantigens and an mRNA product encoding as many as 34 individualized targets is manufactured. The goal is to train the patient's immune system to recognize cancer cells displaying those tumor-specific antigens.

INTerpath-001 enrolled 1,137 patients with completely resected stage IIB-IV melanoma and compared intismeran plus pembrolizumab with pembrolizumab alone. Merck and Moderna reported on August 19, 2026, that the combination produced statistically significant and clinically meaningful improvements in recurrence-free survival and distant-metastasis-free survival. Overall survival remains under evaluation, and the detailed Phase III results have not yet been publicly presented.

Earlier Phase IIB follow-up offers some context. At five years, the companies reported a 49% reduction in the risk of recurrence or death and a 59% reduction in the risk of distant metastasis or death relative to pembrolizumab alone. The larger Phase III readout materially strengthens the evidence that patient-specific antigen selection can produce meaningful treatment.

BioNTech's colorectal result, announced nine days later, showed how much remains unsettled. Its Phase II autogene cevumeran study tested individualized mRNA therapy as monotherapy in ctDNA-positive patients after colorectal-cancer surgery. After an earlier futility boundary was crossed, the independent monitoring board recommended termination because continued treatment was unlikely to alter efficacy and because the data showed a numerical imbalance in overall survival. No new safety signal was identified.

Melanoma and colorectal cancer are biologically different, checkpoint-inhibitor combinations matter, and antigen selection remains difficult. Patient-specific pharmaceutical manufacturing has still reached large randomized testing and is producing meaningful successes and failures.

That creates a practical problem for regulators and insurers. Sequencing, computational target selection, manufacturing and quality control begin to function as parts of one therapeutic platform even though the biological content of the finished product changes with every patient.

Antibodies Become Programmable Therapeutic Systems

Monoclonal antibodies have expanded well beyond receptor blocking or immune marking. Antibody-drug conjugates use an antibody to carry a cytotoxic payload toward cells expressing a target. Bispecific antibodies can use two binding domains to connect a cancer-cell antigen with CD3 on a T cell, physically recruiting the immune system toward the malignancy.

Radioligand therapy applies similar targeting logic with a radioactive payload. In PSMA-positive prostate cancer, a targeting molecule can bind prostate-specific membrane antigen while carrying a radioactive isotope toward cells identified through molecular imaging.

Diagnosis, patient selection and treatment become tightly linked. A growing number of patients can be divided into biologically meaningful subsets and matched with therapies designed for those subsets, while the value of the therapy depends on timely sequencing, pathology, imaging, manufacturing and specialist interpretation.

A system that processes every one of those steps as an unrelated administrative event can become part of the disease clock.

III. Obesity and the Economics of Upstream Treatment

Obesity makes the political consequences of biomedical progress unusually visible because the potential treatment population is enormous. NHANES data for August 2021 through August 2023 found obesity in 40.3% of U.S. adults and severe obesity in approximately 9.4%; another 32.1% of adults were overweight. Prevalence reached 46.4% among adults aged 40 to 59, when diabetes, cardiovascular disease and other metabolic complications begin to accumulate.

CDC's 2026 National Diabetes Statistics Report estimates that 40.1 million Americans had diabetes in 2023, including approximately 11 million adults whose disease was undiagnosed, while 115.2 million adults had prediabetes. CDC estimates diabetes cost the United States \$640 billion in 2021 and accounted for about 25% of health care spending.

Against that background, the GLP-1 revolution is larger than a market for cosmetic weight loss. It creates the possibility of pharmacologically intervening against a major upstream driver of diseases now financed through separate specialties, procedures and lifelong medications.

From Semaglutide to Retatrutide

STEP 1 randomized 1,961 adults without diabetes to semaglutide 2.4 milligrams or placebo alongside lifestyle intervention. After 68 weeks, mean body-weight change was -14.9% with semaglutide compared with -2.4% under placebo.

Tirzepatide subsequently combined GIP and GLP-1 receptor agonism and pushed mean reductions beyond 20% at the highest dose in SURMOUNT-1.

Retatrutide adds glucagon-receptor agonism to those two mechanisms. In its peer-reviewed Phase II trial, the 12-milligram group experienced 24.2% mean weight reduction after 48 weeks compared with 2.1% for placebo, and the weight-loss curve had not clearly plateaued by the end of treatment.

Lilly reported in May 2026 that Phase III TRIUMPH-1 participants receiving 12 milligrams of the investigational drug lost an average 70.3 pounds, or 28.3% of baseline body weight, over 80 weeks; 45.3% lost at least 30%. A 104-week extension among participants with baseline BMI of at least 35 produced 30.3% average weight reduction. At four milligrams, mean loss remained 19%, leaving room for developers to compete across efficacy, tolerability and cost rather than simply chasing the largest possible endpoint.

The Phase III program has also begun reporting disease effects. Lilly has reported reductions in knee-osteoarthritis pain of as much as 73.1%, reductions in obstructive-sleep-apnea severity of as much as 60.6% and A1C reductions of as much as two percentage points in patients with type 2 diabetes. Those remain company-reported pivotal results pending full peer-reviewed publication.

Semaglutide already has hard-outcome evidence. SELECT enrolled 17,604 patients with cardiovascular disease and overweight or obesity but without diabetes and demonstrated a 20% relative reduction in the composite of cardiovascular death, nonfatal myocardial infarction and nonfatal stroke.

Adverse effects, discontinuation and weight regain after stopping therapy all matter to real-world durability. Even with those limitations, a class producing mean weight reductions from roughly 15% to nearly 30% in controlled trials has changed the plausible ceiling of metabolic medicine.

Oral Therapy and the Problem of Scale

Injectable biologic peptides require more complicated manufacturing and distribution than ordinary tablets, and a treatment category measured in tens of millions of potential patients cannot be organized around maximum efficacy alone.

Orforglipron matters for that reason. The small-molecule, nonpeptide oral GLP-1 receptor agonist was evaluated in a Phase III trial of 3,127 adults with obesity. After 72 weeks, mean weight reduction reached 11.1% at the highest studied dose compared with 2.1% under placebo, while 35.9% of patients at the highest dose lost at least 15% of their starting weight. FDA approved the drug as Foundayo on April 1, 2026.

An oral therapy that produces less weight loss than retatrutide can still be disruptive if it is cheaper to manufacture, easier to distribute or acceptable to more patients. High-efficacy injectables can serve one part of the market while oral therapies, lower-dose regimens and future combinations compete elsewhere.

Clinical Value Does Not Automatically Produce Budget Savings

CBO's 2024 analysis of hypothetical broad Medicare anti-obesity-drug coverage estimated \$38.8 billion in federal drug spending between 2026 and 2034 and \$3.4 billion in offsetting healthcare savings, producing approximately \$35.5 billion in net additional federal expenditure. Average drug cost per user declined across the period in CBO's model while medical savings per user rose, yet CBO still projected net Medicare costs through 2044.

Prevention creates this mismatch routinely. The payer buys the medicine today; the avoided heart attack, knee replacement, dialysis episode or diabetes complication may occur years later after the patient has changed employers, moved to another insurer or entered Medicare.

Drug prices also move. The Medicare Drug Price Negotiation Program established a 2027 negotiated price of \$274 for a 30-day supply of selected semaglutide products Ozempic, Rybelsus and Wegovy, compared with a 2024 list price of \$959. Those products accounted for more than \$15 billion in gross Part D spending and were used by roughly 2.28 million beneficiaries in 2024. Rebates and other concessions complicate comparisons with historic net prices, but purchasing power, competition and new products can materially change the economics.

Medicare has already built a temporary structure around the problem. The GLP-1 Bridge operates from July 2026 through the end of 2027 and provides qualifying Part D beneficiaries certain weight-management drugs for a \$50 monthly copayment. Participating products are provided at a \$245 net monthly price to the program, with a central processor handling prior authorization, claims adjudication and pharmacy payment outside ordinary Part D.

The temporary benefit exists because the permanent benefit was not built for the medicine now arriving.

IV. Myostatin and the Next Metabolic Frontier

Aggressive weight loss includes some loss of lean tissue, a concern that matters most for older adults at risk of sarcopenia or frailty. Dieting and bariatric surgery also reduce lean mass, and muscle quality and function matter more clinically than mass alone. Developers nevertheless have an obvious opportunity if they can preserve muscle while metabolic drugs drive fat loss.

Apitegromab targets myostatin, a protein involved in limiting skeletal-muscle growth. In the 102-patient randomized EMBRAZE Phase II trial, participants receiving tirzepatide plus placebo lost 12.5 kilograms overall, including 3.5 kilograms of lean tissue. Participants receiving tirzepatide plus apitegromab lost 11.2 kilograms overall but only 1.6 kilograms of lean tissue, representing 54.9% relative preservation of lean mass. Fat accounted for approximately 85% of weight lost in the apitegromab group compared with roughly 70% in the control group.

The study was too small and too short to establish improved strength, mobility or long-term health outcomes. It still demonstrated that one pathway can drive substantial weight loss while another changes the composition of that loss.

Regeneron's COURAGE program provides an independent test of the same idea. The Phase II trial enrolled 1,005 participants and is evaluating semaglutide combined with trevogrumab, an anti-myostatin antibody, with or without garetosmab, which targets activin A. Company-reported 26-week results indicated that trevogrumab prevented approximately half of the lean-mass loss associated with semaglutide while maintaining weight reduction.

The same biology could matter beyond obesity. Age-related muscle loss contributes to falls, disability, frailty and loss of independence. If developers can preserve or improve muscle while patients lose fat, metabolic treatment begins to overlap with healthy aging.

V. Cell and Gene Therapy and the Shift Toward Biological Correction

Casgevy treats severe sickle cell disease by editing a patient's own hematopoietic stem cells outside the body. CRISPR/Cas9 disrupts regulation of fetal hemoglobin, allowing the edited cells to generate red blood cells that are less likely to sickle. FDA approved the therapy in 2023 and expanded its indication in July 2026 to patients age two and older with sickle cell disease with recurrent vaso-occlusive crises or transfusion-dependent beta thalassemia.

The treatment pathway is much harder than the biological description makes it sound. Patients undergo stem-cell collection, specialized manufacturing, myeloablative conditioning that destroys existing bone marrow cells, reinfusion and prolonged follow-up. Fertility risks, treatment-related complications, long-term surveillance and specialized clinical capacity all matter.

Approval proves that gene editing has become medicine. Delivery remains an industrial and clinical operation.

Casgevy and Lyfgenia entered the U.S. market at list prices of approximately \$2.2 million and \$3.1 million per patient, respectively. Between 50% and 60% of people living with sickle cell disease have Medicaid coverage, leaving state programs with annual budgets disproportionately responsible for therapies whose value may extend over decades.

The financing problem was apparent before either therapy reached the market. In February 2023, CMS selected the Medicaid-focused Cell and Gene Therapy Access Model for testing, with sickle cell disease identified as a potential initial focus. The model now negotiates common outcomes-based arrangements for participating Medicaid programs rather than forcing every state to construct its own performance contract. Thirty-two states, the District of Columbia and Puerto Rico are participating, together representing approximately 84% of Medicaid beneficiaries with sickle cell disease. Manufacturers provide additional rebates when agreed outcomes are not achieved, while CMS supports standardized data collection and implementation.

A one-time \$2 million treatment is not automatically less economical than a chronic therapy costing \$100,000 every year for decades. Annual insurance accounting is simply more familiar with the recurring expense.

Autoimmune Disease and Immune Reset

CAR-T research in autoimmune disease is testing whether a dysfunctional immune system can be depleted deeply enough to reconstitute without recreating the disease.

A 2024 New England Journal of Medicine series followed 15 patients with severe refractory autoimmune disease after CD19 CAR-T treatment; eight had systemic lupus erythematosus, three had idiopathic inflammatory myositis and four had systemic sclerosis. Patients experienced remission or substantial clinical improvement, immunosuppressive therapies were discontinued, and investigators documented B-cell depletion followed by reconstitution.

The evidence remains preliminary. Fifteen highly selected patients without a randomized control group cannot establish comparative effectiveness, long-term safety or durability across broader populations. CAR-T also carries real risk and requires complex manufacturing.

The research question has still changed. Investigators are testing whether selected severe diseases can sometimes be treated episodically at a deeper biological level rather than suppressed indefinitely.

Researchers have also begun moving part of the manufacturing process inside the body. A 2025 New England Journal of Medicine first-in-human report described in-vivo generation of CD19 CAR-T cells using lipid nanoparticles in patients with refractory systemic lupus erythematosus. The treatment produced B-cell depletion and reduced disease activity without the conventional process of collecting, engineering and reinfusing a patient's T cells; the report identified no major toxic effects.

The evidence is extremely early. Reproducing CAR-T biology without ex-vivo cell manufacturing would nevertheless change cost, logistics and scale.

A therapy infused every few weeks indefinitely creates a different economic system from a treatment capable of producing durable remission. If immune resetting proves possible in even a subset of diseases, recurring autoimmune drug franchises, annual insurance spending and provider systems organized around chronic management will all feel the change.

VI. AI, AlphaFold and the Acceleration of Discovery

Drug development will continue to answer to biology. Toxicology, manufacturing, dose selection, human heterogeneity, clinical trials and biological failure do not disappear because a model proposes a promising molecule.

Computational systems can change the search that happens before those bottlenecks. AlphaFold and related structural systems have dramatically expanded access to predicted protein structures and interactions. Machine-learning systems can rank targets, model binding, optimize molecular structures and identify relationships across enormous genomic or phenotypic datasets.

The gain comes from discarding more bad possibilities before expensive experiments begin.

FDA's Center for Drug Evaluation and Research had experience with more than 500 submissions containing AI components from 2016 through 2023 and has developed a risk-based framework for deciding whether models used in regulatory work are credible for their intended context.

Rentosertib provides a useful human marker because both its TNIK target and small-molecule design were developed through an AI-directed platform. In a randomized Phase IIa study involving 71 patients with idiopathic pulmonary fibrosis, the highest-dose group experienced a mean 98.4-milliliter increase in forced vital capacity after 12 weeks compared with a 20.3-milliliter decrease under placebo.

The study was primarily designed for safety. Each treatment group contained fewer than 20 patients, and treatment-related adverse events increased with dose. It was far too small and short to establish clinical efficacy.

The program has since advanced. The first patient in the Phase III GENESIS-IPF-3 trial was dosed on September 10, 2026.

AI-directed target discovery and molecule design have therefore produced a candidate that progressed from a preliminary randomized human signal into late-stage clinical testing while remaining subject to ordinary pharmaceutical uncertainty.

If computational discovery increases the throughput of plausible targets and molecules, the pressure moves into trial recruitment, regulatory review, manufacturing, diagnostics, payer assessment and clinical capacity. Faster discovery has limited value when every institution after discovery still moves at the old speed.

VII. The Peptide Gray Market and the Demand for Enhancement

The contemporary peptide market sits outside the evidentiary core of this report, but it reveals demand that formal medicine cannot simply wish away. Consumers are already seeking pharmacological interventions for injury recovery, body composition, metabolic performance, growth-hormone signaling and longevity before many of the compounds involved have completed conventional development.

BPC-157, TB-500, CJC-1295, ipamorelin, MOTS-C and injectable GHK-Cu are among the products most visibly marketed or discussed in American online and wellness communities. BPC-157 is commonly promoted for gastrointestinal or

musculoskeletal recovery; TB-500 for wound healing and tissue repair; CJC-1295 and ipamorelin for growth-hormone signaling; MOTS-C around metabolism and longevity; and GHK-Cu for regenerative or cosmetic purposes.

No high-quality national surveillance system measures gray-market peptide use, so visibility on social media, in bodybuilding communities or in commercial search results cannot support claims about how many Americans are taking any particular compound.

FDA's regulatory attention shows that the market is no longer obscure. Its Pharmacy Compounding Advisory Committee devoted substantial attention in July 2026 to BPC-157, KPV, TB-500 and MOTS-C, evaluating proposed compounded uses that included ulcerative colitis, wound healing, inflammatory conditions, obesity and osteoporosis. Separate FDA safety reviews describe absent or limited human safety information for several commonly marketed peptides and identify concerns involving immunogenicity, impurities, active-ingredient characterization and, for some compounds, reported serious adverse events.

BPC-157 illustrates the evidence gap. The compound has extensive preclinical literature and substantial anecdotal interest, yet FDA continues to state that it lacks enough information to determine whether compounded BPC-157 can be safely administered through proposed routes. Similar gaps surround injectable MOTS-C and GHK-Cu, while CJC-1295 and ipamorelin present their own safety and characterization concerns.

People want pharmacological tools for recovery, muscle retention, metabolism and aging, and some will accept considerable uncertainty to obtain them. Restricting supply without producing evidence can push that demand toward less controlled sellers; tolerating unsupported claims leaves consumers in uncontrolled experimentation from which medicine learns almost nothing.

Enforcement and research can proceed at the same time. FDA can police contamination, manufacturing and deceptive claims while NIH, academic centers and legitimate developers study compounds attracting sustained real-world use. Standardized toxicology, pharmacokinetic studies and controlled human trials can eliminate ineffective products faster and create a legitimate pathway for any that prove useful.

Demand is not evidence. It can still tell researchers where people are already conducting experiments on themselves.

VIII. Access Is Becoming a Binding Constraint

A treatment has little practical value to an eligible patient who cannot receive it. More powerful therapies are exposing weaknesses in an access system built around formularies, annual contracts, prior authorization, step therapy and relatively standardized pharmaceuticals.

A KFF analysis of CMS reporting data found that Medicare Advantage insurers made approximately 52.8 million prior-authorization determinations in 2024 and denied roughly 4.1 million requests in whole or in part. Only 11.5% of those denials were appealed, but 80.7% of appealed denials were ultimately overturned.

Additional documentation or changed circumstances can affect an appeal, so that reversal rate does not prove that four out of five initial decisions were irrational. It does show a large administrative process in which the final answer often differs from the first one.

Physicians describe the burden from the other side. The 2025 AMA Prior Authorization Physician Survey found that practices completed approximately 40 prior authorizations per physician each week and devoted roughly 13 hours of physician and staff labor to them. Forty percent employed staff devoted exclusively to authorization work, while 95% of physicians said the process delayed necessary treatment and 79% said patients sometimes abandoned treatment because of the burden.

CMS has already begun rebuilding the machinery. Its 2024 prior-authorization rule imposed new electronic and timing requirements for non-drug services. In April 2026, the agency proposed a separate pharmaceutical rule extending electronic authorization standards to drugs and, in several federal programs, requiring decisions within 24 hours for urgent requests and 72 hours for standard requests beginning in October 2027.

American laboratories can now sequence a tumor and manufacture an individualized mRNA therapy around its mutations while the government is still establishing uniform electronic rules for how an insurance plan says yes or no to the prescription.

Step Therapy and Precision Medicine

Step therapy makes economic sense when several treatments are genuinely substitutable and a lower-cost option produces comparable outcomes. Precision medicine erodes that assumption when a validated biomarker identifies the subgroup most likely to respond to a particular therapy.

Molecular pathology can identify receptors, mutations, HLA types and antigens that determine therapeutic eligibility. A utilization-management system designed for interchangeable mass-market drugs can otherwise require a patient to fail an alternative after the diagnostic evidence has already identified the more appropriate treatment.

Companion diagnostics belong in the same clinical pathway. Coverage for a drug that requires genomic sequencing, molecular imaging or antigen testing is of limited value if the payer separately denies the test establishing eligibility.

The more precise the medicine becomes, the stranger it looks to finance the diagnostic and the treatment as unrelated decisions.

PBM Concentration

Pharmacy benefit managers add another layer of gatekeeping. FTC reported that the three largest PBMs processed nearly 80% of approximately 6.6 billion prescriptions dispensed by American pharmacies in 2023, while the six largest processed more than 90%. Pharmacies affiliated with the three largest PBMs accounted for nearly 70% of specialty-drug revenue.

PBMs influence formulary placement, pharmacy networks, utilization management and patient cost at the same time that specialty pharmaceuticals are becoming more heterogeneous. A formulary model developed around competing cholesterol medicines or proton-pump inhibitors fits awkwardly with individualized oncology, cell therapy and rare-disease products that may have no straightforward substitute.

Comparable public metrics would make that private decision-making easier to evaluate. Denial rates, appeal rates, reversal rates, decision times and the evidence supporting coverage restrictions should be visible by therapeutic category.

Without them, policymakers are trying to distinguish legitimate utilization management from administrative friction while looking through a black box.

IX. The Political Economy of Therapeutic Abundance

Drug manufacturers, insurers, employers and public programs often operate on different clocks.

A manufacturer developing a durable treatment can value years of expected clinical benefit. The insurer paying the claim may cover the patient for only a year or two. If the therapy prevents complications a decade later, another insurer, another employer or Medicare may capture much of the benefit. A state Medicaid program can receive a multimillion-dollar gene-therapy bill this year even when the treatment may save public money over a much longer period.

The obesity market produces the same mismatch at population scale. Large employers increasingly acknowledge the effectiveness of GLP-1 therapy, yet Mercer found that 6% dropped weight-loss GLP-1 coverage in 2026 and another 5% were considering or planning to do so for 2027; 27% tightened or planned to tighten utilization controls.

The arithmetic becomes difficult when a highly effective chronic medicine is clinically relevant to a large share of the workforce.

One-time treatments create the opposite cash-flow pattern. The payer absorbs a concentrated claim while the patient receives future value, which can make chronic recurring expenditure easier to budget than a durable intervention even when the lifetime economics favor the latter.

CMS's cell-and-gene model attempts to bridge part of that gap by tying some payment to outcomes and pooling negotiation across states. Outcomes-based rebates, reinsurance, risk corridors and carefully designed multi-year financing may become more important as durable therapies expand. Every one of those approaches eventually encounters the same practical problem: somebody has to track outcomes across years and across changes in insurance coverage.

Pricing Remains Part of the Problem

Access cannot mean accepting any launch price a manufacturer chooses. If every incremental innovation is priced as though it were curative, therapeutic progress becomes fiscally impossible and high prices themselves ration access.

The metabolic market shows how competition can work on more than one dimension. Semaglutide, tirzepatide, oral orforglipron and investigational retatrutide compete across efficacy, convenience and delivery, while federal negotiation has materially lowered the future Medicare price for selected semaglutide products.

Clinically important innovation needs returns large enough to support expensive and risky research. Successful therapeutic classes also need competitors, alternative delivery methods and manufacturing improvements that push prices downward as the class matures.

Political Expectations Will Change

Patients know more about which medicines exist before they enter the examination room. Social media, direct-to-consumer telehealth and mainstream news have made GLP-1 coverage decisions unusually visible, and gene or cancer therapies raise the stakes further when delay can affect survival or eligibility for treatment.

An administrative delay means something different when cancer is progressing or a patient must remain healthy enough to complete a cellular-treatment process.

Insurers and PBMs will face growing pressure to explain why a restriction is clinically justified and how quickly it can be resolved. Manufacturers whose prices appear disconnected from measurable benefit will face the same pressure from employers and public payers.

Better medicine makes every gatekeeper easier to see.

X. Delivery Capacity Is Part of Access

Insurance approval does not create a manufacturing slot, a trained treatment team or a radiopharmaceutical supply chain.

CMS's description of the sickle-cell gene-therapy process spans months of outpatient and inpatient care, including identification, specialist evaluation, cell collection, manufacturing, conditioning and infusion. Some patients also need fertility-preservation services because the myeloablative treatment used before gene therapy can cause infertility.

CAR-T requires leukapheresis, cellular manufacturing and clinicians capable of managing cytokine-release syndrome and neurologic toxicity. Personalized cancer treatment depends on tissue acquisition, sequencing, bioinformatics and rapid individualized production. Radioligand therapy requires reliable isotope production, specialized transportation, nuclear-medicine personnel and facilities licensed to handle radioactive materials.

That infrastructure is unevenly distributed. Advanced-treatment centers remain concentrated around large academic or specialized hospitals, leaving rural patients and people in smaller metropolitan areas with travel, lodging, time-off-work and caregiver burdens even when insurance nominally covers the therapy.

Therapeutic capacity includes molecular pathologists, genetic counselors, cellular-manufacturing staff, nuclear pharmacists, infusion nurses, transplant teams, specialized pharmacists, data systems and logistics networks.

A country can own the intellectual property for extraordinary medicine and still ration access through shortages of trained people, production capacity and treatment slots.

XI. Regulation for an Era of Individualized Medicine

The twentieth-century regulatory model developed largely around standardized pharmaceutical products manufactured repeatedly for large populations. Much of medicine still fits that model. Individualized therapies complicate the assumption that every unit of the product is biologically identical from patient to patient.

FDA has begun adapting. In February 2026, the agency issued draft guidance describing a plausible-mechanism framework for individualized treatments targeting genetic diseases with known biological causes. The framework specifically contemplates genome editing and RNA-based therapies in populations so small that conventional randomized trials may be impractical while continuing to require evidence supporting safety, effectiveness and manufacturing quality.

Individualized cancer platforms create a related problem. When every cancer vaccine encodes different tumor mutations, regulators need confidence in the sequencing, computational selection, manufacturing and quality-control process even though the patient-specific payload changes.

Treating every customized sequence as though it were an unrelated drug would eventually make individualized treatment impossible to scale. Treating the platform as automatically trustworthy would create a different problem.

Durable genetic and cellular therapies can create risks that emerge years after administration, so flexibility before approval has to be paired with strong postmarket data, registries and longitudinal surveillance.

AI introduces another version of the same issue because a model used for target identification, trial selection or manufacturing can change over time. FDA's emerging risk-based approach should focus scrutiny on how consequential the model is to a regulatory decision, the quality of evidence supporting its performance and whether changes in data or implementation alter reliability.

The useful standard is simple: scrutinize genuine uncertainty aggressively and stop making well-characterized platforms repeat procedures that add little information.

XII. A Policy Framework for Therapeutic Abundance

More capable medicine requires a system that can distinguish quickly between uncertainty and a patient who plainly satisfies established clinical criteria.

Prior authorization is an obvious place to start. Diagnosis, biomarker results, previous treatment and laboratory values increasingly exist as structured information inside the medical record. When those facts satisfy published coverage criteria, much of the authorization process can be automated. Human review remains useful when clinical judgment is actually required; using it as a universal checkpoint consumes skilled labor and adds delay.

Decision times should follow clinical urgency. CMS's proposed 24-hour and 72-hour standards for several classes of drug authorization provide a useful baseline, particularly where delay can affect disease progression or eligibility for treatment.

Step-therapy rules also need to recognize biomarkers. Insurers have a legitimate reason to require a less expensive treatment when the alternatives are clinically equivalent. Validated molecular evidence can eliminate that equivalence for a particular patient. A receptor, mutation, HLA type or antigen should be able to satisfy the coverage rule before a patient is required to prove the same point by failing another treatment.

Genomic sequencing, molecular imaging and biomarker testing should be financed as part of the therapeutic pathway when they determine whether an expensive drug should be used. Denying the test can block the effective treatment and can also encourage inefficient prescribing by removing the evidence needed to select patients well.

High-cost durable therapies need financing mechanisms that can see beyond one insurance year. The CMS Cell and Gene Therapy Access Model provides one template through outcomes-linked rebates. Reinsurance, broader risk pooling and portable multi-year financing deserve further development where the clinical value of a therapy may extend across decades.

Treatment capacity belongs in the same conversation. Advanced biomanufacturing, isotope supply chains, workforce training, cellular-therapy centers and genomic infrastructure become healthcare infrastructure once an approved treatment depends on them. Approving a therapy without expanding the ability to administer it simply moves rationing from the insurance desk to the waiting list.

FDA should continue developing platform-based regulatory approaches for individualized and ultra-rare therapies while strengthening postmarket surveillance and maintaining clear standards for meaningful clinical benefit. Flexibility is useful when it removes redundant evidence requirements; it becomes dangerous when it lowers the threshold for showing that a treatment works.

The peptide gray market requires a different approach. Products marketed for human use under nominal research labels should face enforcement when purity, sterility or claims are deceptive. Compounds that attract sustained real-world use can simultaneously become candidates for accelerated independent study. The goal should be to replace uncontrolled experimentation with evidence.

PBM and insurer transparency should improve as well. Public reporting of denial rates, decision times, appeal rates, reversal rates and major coverage criteria by therapeutic category would give policymakers a better basis for evaluating whether utilization management tracks clinical evidence.

Competition remains essential. Oral alternatives, next-generation products, biosimilars, generics where scientifically possible, manufacturing innovation and public purchasing power can lower cost as successful therapeutic classes mature. The system has to reward important innovation strongly enough to sustain research without turning every successful product into a permanent tollbooth.

Longer time horizons tie the pieces together. CBO scoring rules, employer benefits and annual insurance contracts will remain, but policymakers need to distinguish a treatment that increases this year's pharmacy spending from one that increases lifetime healthcare costs.

Medicine is beginning to operate across biological time horizons that do not fit neatly inside an insurance year.

Conclusion

The argument for therapeutic abundance no longer rests primarily on what might work.

Semaglutide has reduced major cardiovascular events in a randomized trial involving 17,604 adults. Investigational retatrutide has produced Phase III weight reductions approaching 30%. An FDA-approved oral small-molecule GLP-1 has opened a more scalable delivery model. CRISPR-edited stem cells are approved medicine. Individualized mRNA cancer treatment has produced a positive Phase III readout. Patient-derived immune cells and engineered T-cell receptors are approved for solid tumors. Randomized studies are testing whether muscle can be preserved during pharmacological weight loss, while early CAR-T research is exploring whether selected autoimmune diseases can move from lifelong suppression toward durable remission.

Some of these efforts will fail. BioNTech's colorectal-cancer result in August 2026 already did. Larger controlled trials may weaken the early autoimmune results. Gene therapies will remain applicable to small populations in many diseases, and metabolic drugs may require sustained treatment. Therapeutic abundance does not mean every experiment succeeds.

It means enough different biological systems are becoming clinically tractable that failure is increasingly happening inside a much larger field of things medicine can seriously attempt.

The harder problem is moving downstream.

Annual insurance contracts finance treatments whose value may extend for decades. Pharmacy benefits, medical benefits and diagnostic coverage remain separated even when one treatment pathway requires all three. Utilization-management systems designed for broadly interchangeable drugs are being applied to therapies selected by molecular biology. Specialized manufacturing and treatment capacity remain concentrated in a relatively small number of places.

Policymakers are already improvising around those mismatches. Medicare built the GLP-1 Bridge outside ordinary Part D. CMS's Cell and Gene Therapy Access Model is creating a national outcomes-based framework for sickle-cell gene therapy because state Medicaid programs are poorly positioned to negotiate and absorb multimillion-dollar one-time treatments independently.

The next generation of therapies will put more pressure on the same weak points. Cell and gene therapies require specialized hospitals and months-long treatment pathways. Radiopharmaceutical growth depends on isotope supply chains and nuclear-medicine capacity. Individualized oncology depends on sequencing, pathology, data systems and manufacturing that can move on the timetable of a patient's disease.

A patient can have the right diagnosis, an approved therapy and insurance coverage and still discover that there is no treatment slot, no covered diagnostic, no manufacturing capacity or no timely authorization.

For decades, healthcare policy could assume that medicine itself imposed many of the hardest limits. In important areas, those limits are beginning to move.

The institutions behind the medicine now have to move with them.