Biotech and Pharma Forecasting

Where Forecasts Break Down, Why It Happens, and How to Build More Resilient Forecasts

This article draws on peer-reviewed literature and official FDA and EMA materials to examine where forecasting models most often fail in pharmaceutical and biotechnology settings. It covers commercial launch forecasting, in-line demand forecasting, clinical development forecasting, portfolio decision support and supply planning.

DSX

Executive Summary

Forecasting failures in pharmaceutical and biotechnology companies are often treated as technical or algorithmic problems. The evidence suggests that many forecast errors arise from the design of the model and the assumptions within it, not from the mathematics.

Formula and coding errors do occur, but the larger failures usually happen when the forecast does not reflect the real system it is meant to describe: the treatable patient population, the treatment journey, the launch curve, the access environment, the clinical pathway, or the supply system required to deliver product. Even a sophisticated method will fail if the model is built on the wrong assumptions.

Across commercial, clinical, and operational settings, seven common forecast breakdowns appear repeatedly:

1	**Patient Population** Misunderstanding the patient population and treatment journey.
2	**Launch Analogues** Using inappropriate launch analogues and unrealistic time-to-peak assumptions.
3	**Access & Policy** Underestimating how much access, reimbursement, competition, and policy can change over the forecast horizon.
4	**Clinical Uncertainty** Treating clinical enrolment, timelines, and probability of success as fixed rather than uncertain.
5	**Demand–Supply Gap** Disconnecting demand forecasts from manufacturing and supply-chain constraints.
6	**Rare Disease** Misreading the low-volume, high-complexity economics of rare disease and specialty markets.
7	**Override Governance** Allowing anecdotal or unsupported overrides without proper governance or evidence review.

These are not rare cases. They are the common ways forecasts break down, with direct consequences for launch planning, investment decisions, and operations.

Better forecasting rarely starts with a more complex model. It starts with better forecast design: patient-based models where treatment flow matters, scenario-based approaches where market conditions may shift, probabilistic methods where uncertainty cannot be reduced, and close links between demand and supply where operational limits affect what can actually be delivered.

Best practice is not a single-point forecast, but a forecasting approach that makes assumptions explicit, separates evidence from judgement, and is updated quickly when clinical, regulatory or market events change the outlook.

Key Conclusions

- Forecast failures are more often caused by flawed assumptions about patients, access, competition, or development than by weaknesses in modelling technique.
- In oncology, rare disease, and other specialty markets, patient-level and treatment-journey logic matter more than top-down market sizing.
- Single-point forecasts are inadequate for launch, reimbursement, clinical development, and supply decisions exposed to discontinuous events.
- Human judgement remains necessary, but undocumented overrides and consensus adjustments often reduce forecast quality.
- Strong forecasting capability links commercial, clinical, finance, and technical operations in a single evidence-based process.

Why Forecasting Is Challenging in Pharmaceutical and Biotechnology Markets

Forecasting in pharma and biotech is materially different from forecasting in most consumer, industrial, or retail markets. Demand is not simply driven by price, promotion, and seasonality. It is the outcome of a chain of events: disease burden, diagnosis, testing, medical eligibility, referrals, system capacity, reimbursement, physician adoption, patient persistence, and supply. In many therapy areas, forecasts fail not because the arithmetic is weak, but because one or more parts of that chain have been omitted, oversimplified, or assumed to be stable when they are not.

The challenge is compounded by the way forecasts are used. They support very different decisions: asset valuation, portfolio investment, clinical planning, launch preparation, manufacturing capacity, investor communication, and product supply. Those decisions do not all require the same forecast. Yet in practice, one model is often expected to do all of them. The result is a hybrid of commercial expectation, internal advocacy and operational planning, with less clarity about what the forecast is actually trying to predict.

A further complication

Some of the most important drivers are discrete events rather than gradual trends. A pivotal trial readout, competitor approval, reimbursement decision, guideline change, manufacturing deviation, or facility shutdown can all alter the market quickly.

In these settings, the right question is often not **"What is the number?"** but **"What are the plausible paths, what could cause the path to change, and what decision should we make under each one?"**

EMA guidance on demand forecasting makes this point clearly: forecasts should use the best available information, cover at least six months when manufacturing planning is involved, and be updated regularly because circumstances can change quickly.

> A strong pharma forecast is therefore not just a single number. It is a structured view of how patients, evidence, access and supply are likely to evolve.

Where Forecasts Commonly Break Down

Failure Mode	Typical Symptom	Business Consequence	Preferred Design Response
Patient population and flow error	Peak sales look attractive, but diagnosed, tested, referred, or treatable patient volume does not support implied uptake	Inflated launch expectations and poor resource allocation	Patient-based funnel with explicit drop-offs and graded evidence
Launch curve and analogue bias	Revenue phasing is too fast or too slow relative to actual uptake	Distorted NPV, launch investment, and BD valuations	Archetype-specific launch libraries with event-triggered re-estimation
Access and policy-change blindness	Strong approval but weak actual uptake or marked regional divergence	Missed access expectations and delayed ramp	Separate approval from effective access; scenario-test payer and policy shifts
Clinical timeline overconfidence	Enrolment and milestones repeatedly slip	Portfolio and cash-flow distortion	Recruitment scenarios and asset-specific success assumptions
Demand–supply decoupling	Demand assumptions are strong but shipments miss because product is unavailable	Lost revenue, shortages, and reactive firefighting	Constrained-supply forecasting linked to quality and capacity assumptions
Rare disease / advanced-therapy volume distortion	Small assumption changes create very large forecast swings	Misbuilt manufacturing and network design	Patient-level or centre-level simulation; separate prevalent and incident cohorts
Override and governance failure	Consensus forecast changes every cycle with little learning	Bias, wasted effort, and model distrust	Preserve the baseline, log overrides, and review forecast value added

The Most Common Areas Where Forecasting Models Fail

Patient Population and Treatment Journey

Forecasts fail when prevalence, incidence or total addressable market are treated as if they were the same as the number of patients who will actually be treated. In oncology, specialty care, rare disease, and advanced therapies, the treated population is shaped by a series of filters. Patients may remain undiagnosed, may not receive the relevant biomarker test, may fail clinical eligibility criteria, may not reach a qualified centre, may not secure reimbursement in time, or may discontinue earlier than assumed. When a model jumps from epidemiology to revenue without modelling these filters explicitly, it usually overstates demand.

Epidemiology matters and can appear precise, but real patient journeys are often fragmented. Commercial, medical, and market access teams often hold different parts of the picture, and there is a risk that the forecast ends up reflecting whichever version is easiest to model. Yet the direction of the evidence is clear: epidemiology is the starting point, not the forecast itself. The EMA explicitly describes epidemiological modelling as the foundation for medicine demand forecasting and stresses the need to account for uncertainty. Recent patient-based work on cell and gene therapies takes the same approach, combining prevalence, incidence, time to launch, probability of success, eligibility and adoption rather than relying on a single market multiple.

Best practice is to build a transparent patient funnel. At minimum, the model should distinguish prevalence or incidence, diagnosed patients, tested patients where relevant, medically eligible patients, reimbursed or access-eligible patients, and operationally treatable patients. Each transition should have an evidence owner, a refresh cadence, and a confidence grade. In rare diseases and advanced therapies, referral capacity and centre readiness should sit inside the forecast rather than being treated as downstream issues.

Launch Analogues and Time-to-Peak Assumptions

In the launch phase, forecasts are inherently uncertain. The important question is not whether the initial curve will be wrong, but how wrong it is likely to be and how quickly it can be updated as real uptake data emerges.

Many launch models rely on analogue curves built from superficially similar brands, for example those with the same mechanism, therapy area, or price range. This can be useful, but it is also risky. Small differences in access restrictions, route of administration, evidence quality or competitive timing can cause uptake to differ materially from the analogue set.

In other cases, the model does not use a true analogue at all, but instead applies a default time-to-peak assumption because it is convenient for valuation or planning. That may produce a neat model, but it may still be poorly matched to the asset and the market.

Robey and David show that shortening time to peak can significantly change NPV. Regnier and Ridley find that order of entry, promotional share, and time to market all influence peak share. Donoghoe and colleagues show that fewer launches are meeting expectations, and that success has become harder in a more competitive and access-constrained environment.

Generic analogue curves are therefore rarely sufficient. Oncology brands, rare disease biologics, primary care small molecules, and gene therapies cannot be expected to follow the same default uptake logic. The adoption model should separate awareness, willingness to prescribe, access friction, provider or site readiness, and supply availability. Early post-launch learning should focus on the drivers of adoption, not just top-line sales variance, so the curve can be recalibrated quickly in the first 12 to 24 months after launch.

Policy and Access

Forecasts fail when regulatory approval is treated as equivalent to real access, or when access is reduced to a simple delay after launch. In practice, uptake is shaped by payer restrictions, HTA outcomes, pathway decisions, formulary position, prior authorisation, contracting, and ongoing evidence requirements. Competition adds further complexity. Order of entry, label breadth, comparative evidence, and pricing can all change both the speed of uptake and the level of long-term share, often in non-linear ways.

Access is also often added too late, once the patient and demand logic has already been set. That tends to understate regional friction and overstate the speed of uptake. Donoghoe and colleagues describe a launch environment that has become harder, shaped by greater payer scrutiny, pricing pressure, and tighter access controls. EMA guidance makes the same point: forecasts can become outdated quickly, alternative scenarios should be explored, and policy decisions are among the hardest variables to anticipate reliably.

A strong forecast should separate three distinct stages: **regulatory approval**, **reimbursed access**, and **practical access** in the real treatment setting. For more volatile products, monthly review may be more appropriate than quarterly review — not because the forecast is broken, but because the market may have changed.

Trial Enrolment, Milestones and Probability of Success

The same forecasting discipline should be applied to clinical development, which is often treated as more predictable than it is. Deterministic enrolment curves, standard phase-transition assumptions, and fixed probability-of-success inputs often create a false sense of precision around milestone timing, trial duration, and portfolio value.

In practice, clinical timelines are uncertain and should be modelled that way. Schachter and Ramoni note that forecasting late-stage assets remains difficult because early data rarely provide enough certainty to stop future failures without also risking good programmes. Systematic review evidence shows that failure to recruit to time and target is frequent, often because recruitment rates and eligible patient numbers are overstated. Benchmarking of phase III recruitment highlights the cost of operational delay: in a $150 million phase III programme run over three years, around **$1 million per week is at risk if the trial slips for operational reasons.**

The practical response is to make clinical forecasting probabilistic by design. Enrolment should be modelled to reflect uncertainty around site activation, site productivity, screen-failure rates, and competition from other trials. In most cases, the right output is not a single milestone date, but a range of plausible dates linked to clear scenarios and trigger points.

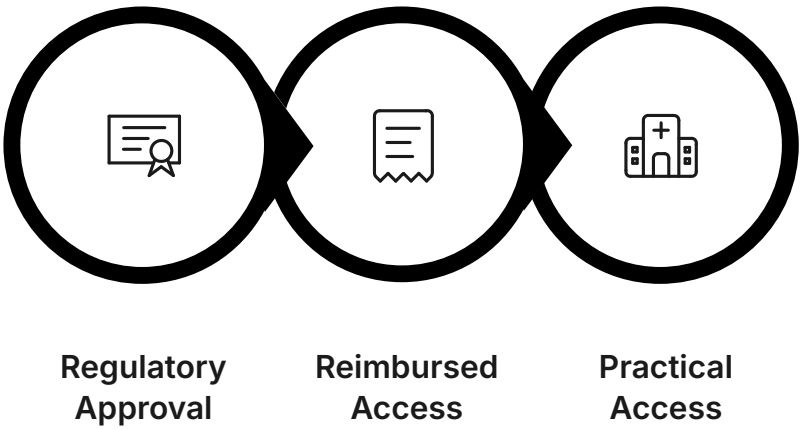

Separating these three stages prevents the common error of treating approval as equivalent to uptake, and makes the model account for the real-world friction that shapes the launch curve.

Demand Forecasting and Supply Chain

Demand models fail when they assume product will be available simply because patient demand exists. This is one of the most common blind spots in demand planning. A commercial team may estimate in-market demand correctly and still miss revenue because of quality issues, capacity constraints or other supply challenges.

Many organisations still maintain separate forecasts for commercial demand and supply planning, but in practice the two cannot be separated. EMA guidance states clearly that forecasting is essential for making the right adjustments in manufacturing and distribution to avoid or mitigate shortages. FDA research on drug shortages identifies three structural drivers: weak incentives to supply less profitable products, limited reward for mature quality systems, and logistical and regulatory barriers that slow recovery after disruption. In its CY 2024 report to Congress, the FDA notes that shortages remain a significant public health challenge.

> Commercial demand, shippable supply and supply available for allocation should be represented separately. Supply forecasts should reflect batch success, release timing, quality risk, sourcing concentration and recovery scenarios. Shortage prevention should be built into forecast design rather than treated as a separate crisis process.

Rare Disease and Specialty Markets

Forecasts in rare disease, specialty, and advanced therapy markets do not behave like conventional chronic therapy models. In these settings, small changes in a single assumption — such as patient identification, referral success, or reimbursement timing — can materially change the forecast. One-time therapies create particularly unusual demand curves because the initial treatment of the prevalent pool produces a front-loaded wave of demand, followed by a steadier state driven mainly by incident patients.

Planning in rare disease and advanced therapy often starts with a headline prevalence figure, which can create a false sense of the addressable patient population. Recent patient-based work in France shows that forecasting advanced therapies requires explicit assumptions about pipeline timing, probability of success, epidemiology, eligibility and adoption. The CAR-T supply-chain literature points to clear limits: processes are difficult to scale and depend on complex logistics, as well as being constrained by specialist-centre capacity and reimbursement complexity.

Forecasts in these markets should separate prevalent and incident cohorts, model patient identification and referral explicitly, and simulate centre-level throughput where capacity is tight.

Governance

Reliable forecasting depends on clear governance around how judgement is used. The issue is not that experts intervene. That is often necessary, because new evidence, access signals and operational events do matter. The issue is that many organisations neither preserve the baseline model output nor record why overrides were made.

Launch pressure, portfolio politics, optimism bias, and consensus culture all encourage manual adjustment. In a 15-year case study of a pharmaceutical company, Fildes and Goodwin found little difference between judgement-based and automatic baseline forecasts, despite the extra effort involved. Truebel and Seidler's review of R&D decision-making also highlights the role of bias among senior decision-makers and the need for systematic mitigation.

Organisations should preserve the model baseline, record every override together with its rationale and supporting evidence, and review those adjustments systematically to determine whether they improve the forecast.

The process should also separate baseline generation from the wider consensus discussion, so that the analytical baseline is not silently rewritten.

Building a Resilient Forecasting Framework

The strongest response to these recurring forecasting challenges is not a single new model. It is a redesign of the forecasting system. The framework below translates those issues into a more resilient forecasting approach.

1

Build one causal architecture with multiple outputs

Commercial, clinical, and supply decisions should share core assumptions where they refer to the same underlying reality, but they should not be forced into a single undifferentiated number. One causal architecture can support multiple decision views: patient demand, reimbursed demand, shipped demand, milestone ranges, and valuation scenarios.

2

Make assumptions explicit and grade the evidence

Every major assumption should have a source, refresh cadence, and confidence grade. This is particularly important for diagnosis rates, testing rates, eligibility, access timing, probability-of-success inputs, and supply-chain assumptions.

3

Use probabilities where uncertainty cannot be removed

Trial enrolment, milestone timing, pivotal success, manufacturing recovery, and advanced-therapy capacity should usually be modelled as distributions or scenario ranges rather than single points.

4

Use event-driven scenarios

Forecasts should specify which external events trigger re-estimation: competitor data, HTA outcomes, label changes, safety events, supply shortages, site openings, plant shutdowns, and policy actions. This allows the forecast to evolve in a controlled way rather than through ad hoc reactions.

5

Match the model to the product archetype

A primary care chronic therapy, an infused oncology medicine, an orphan biologic, and an autologous cell therapy each require different forecast logic. Modality-specific templates reduce the temptation to force every product into the same spreadsheet architecture.

6

Govern overrides and test whether they add value

Human intervention should be deliberate, logged, and reviewable. The organisation should know whether specific types of intervention improve forecast performance, worsen it, or simply consume time.

7

Close the learning loop

An inaccurate forecast should be audited at the level of assumptions, not only at the level of top-line error. The goal is not just to know that the forecast was wrong. It is to understand whether the miss came from patient volume, access timing, launch shape, supply, or managerial intervention.

Forecast Archetype Guide

Forecast Archetype	Primary Engine	Key Uncertainties to Model Explicitly	Common Trap
Primary care / broad chronic therapy	TRx / NRx and channel-based demand model	Generic erosion, contracting, payer mix, adherence	Using patient funnels where prescription and channel data are stronger demand signals
Specialty biologic	Patient funnel with explicit access layer	Referral flow, access timing, site of care, persistence	Treating regulatory approval as equivalent to reimbursement
Precision oncology	Epidemiology-to-treatment funnel, from biomarker to line-of-therapy use	Testing rates, competing regimens, centre adoption, evidence evolution	Treating incidence as if it were treated demand
Rare disease / orphan	Identified-patient model with separate prevalent and incident cohorts	Diagnosis speed, referral capture, access timing, centre capacity	Assuming prevalence converts into actual treated patients
Autologous cell / gene therapy	Centre-level throughput model linked to constrained supply	Work-up, end-to-end turnaround time, manufacturing success, reimbursement timing, scheduling	Modelling monthly demand as if it were a conventional infused product
Development portfolio forecast	Probabilistic milestone and probability-of-success model	Recruitment variability, protocol changes, evidence translation, rare vs non-rare benchmarks	Using blanket historical success rates in place of asset-specific evidence

Implementation Roadmap

A mature forecasting capability is built in stages. The sequence below is designed to show visible progress within one planning cycle while laying the groundwork for a more resilient system.

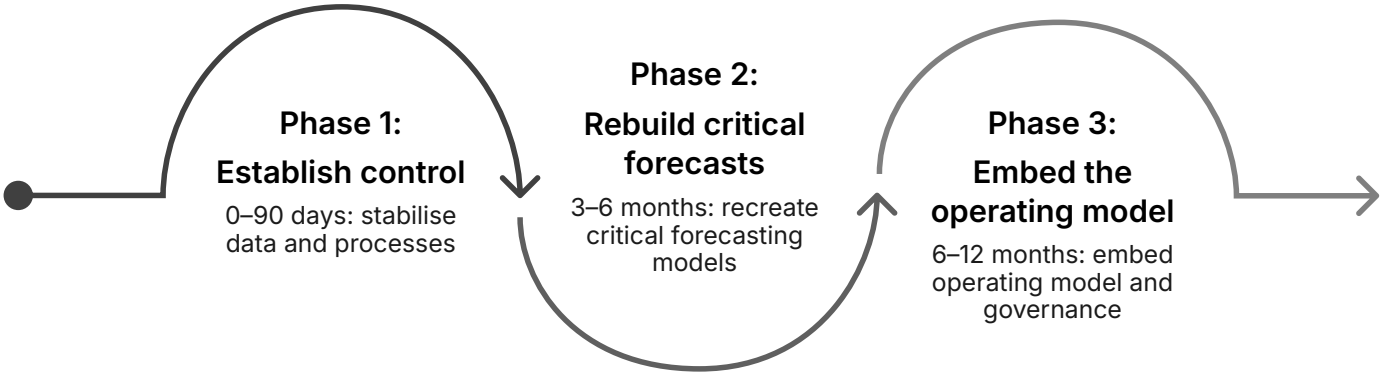

Each phase builds on the last, ensuring that governance and assumption control are in place before more complex model rebuilds begin.

1 Phase 1: Establish control

- Audit the forecast types currently in use: portfolio, development milestones, launch, in-market demand, and supply.
- Identify the highest-impact historical forecast misses and classify them by failure mode rather than by business unit.
- Create a shared assumption register covering patient volume, access, launch timing, clinical milestones, and supply requirements and constraints.
- Preserve baseline model outputs and begin logging overrides with their rationale, magnitude, and evidence source.
- Define decision-specific performance metrics so that launch, supply, and portfolio forecasts are not all judged by a single generic accuracy measure.

2 Phase 2: Rebuild critical forecasts

- Rebuild the highest-priority commercial forecasts using patient-based funnels where relevant.
- Introduce event-triggered scenarios for reimbursement, competitor events, and manufacturing disruptions.
- Pilot probabilistic recruitment and milestone forecasting for the most important development programmes.
- Separate requested demand, constrained supply, and supply available for allocation in operational and launch planning.
- Begin quarterly forecast reviews to test which management overrides improve the process and which should be curtailed.

3 Phase 3: Embed the operating model

- Create a cross-functional forecasting governance forum spanning commercial, medical, market access, finance, development, and technical operations.
- Deploy modality-specific forecasting templates for primary care, specialty, oncology, rare disease, and advanced therapies.
- Connect post-mortems to model redesign so that forecast misses improve the model rather than simply shifting accountability.
- Introduce calibration reviews for probability forecasts, particularly for clinical milestones and launch scenarios.
- Apply the same causal logic across strategy, launch readiness, and supply planning so that inconsistencies between models are visible early.

Conclusion

The main lesson from practice and published research is that forecasting models tend to fail in recognisable ways. They fail when the patient model is too simple, launch curves are too optimistic, access friction is ignored, clinical timelines are treated as fixed, supply constraints are overlooked, and human intervention is not properly governed.

That is as much an opportunity as a warning. Many of the largest forecast errors are not random. They are design failures. Organisations that model patient flow, uncertainty, access and constrained supply explicitly — and govern overrides through evidence and review — can improve forecast quality without waiting for a better algorithm.

Model Patient Flow

Build transparent patient funnels with explicit drop-offs, graded evidence, and clear evidence owners at each transition.

Embrace Uncertainty

Use probabilistic methods for clinical milestones, enrolment, and supply recovery rather than single-point assumptions.

Model Real Access

Separate regulatory approval, reimbursed access, and practical access. Scenario-test payer and policy shifts explicitly.

Connect Demand and Supply

Represent commercial demand, shippable supply, and supply available for allocation as distinct layers in the forecast.

Govern Overrides

Preserve the baseline, log every override with its rationale, and review systematically whether interventions add value.

The goal is not perfect foresight. It is a forecasting system that makes its assumptions explicit, supports decisions, and adapts quickly when conditions change.

DSX

References

The following peer-reviewed publications and official regulatory materials informed this article.

- **1.** Cha M, Rifai B, Sarraf P. Pharmaceutical forecasting: throwing darts? *Nature Reviews Drug Discovery.* 2013;12:737–738. doi:10.1038/nrd4127.
- **2.** Regnier SA, Ridley DB. Forecasting market share in the US pharmaceutical market. *Nature Reviews Drug Discovery.* 2015;14:594–595. doi:10.1038/nrd4697.
- **3.** Robey S, David FS. Drug launch curves in the modern era. *Nature Reviews Drug Discovery.* 2017;16:13–14. doi:10.1038/nrd.2016.236.
- **4.** Donoghoe N, Duane J, Kim J, et al. Pulling away from the pack in drug launches. *Nature Reviews Drug Discovery.* 2017;16:749–750. doi:10.1038/nrd.2017.122.
- **5.** Schachter AD, Ramoni MF. Clinical forecasting in drug development. *Nature Reviews Drug Discovery.* 2007;6:107–108. doi:10.1038/nrd2246.
- **6.** Gkioni E, Rius R, Dodd S, Gamble C. A systematic review describes models for recruitment prediction at the design stage of a clinical trial. *Journal of Clinical Epidemiology.* 2019;115:141–149. doi:10.1016/j.jclinepi.2019.07.002.
- **7.** Basu SB. Benchmarking recruitment rates for phase III trials. *Nature Reviews Drug Discovery.* 2024;23:887–888. doi:10.1038/d41573-024-00183-4.
- **8.** Chen C, Zhou X, Lavezzi SM, Arshad U, Sharma R. Concept and application of the probability of pharmacological success (PoPS) as a decision tool in drug development. *Journal of Translational Medicine.* 2023;21:17. doi:10.1186/s12967-022-03849-y.
- **9.** Hampson LV, Holzhauer B, Bornkamp B, et al. A new comprehensive approach to assess the probability of success of development programs before pivotal trials. *Clinical Pharmacology & Therapeutics.* 2022;111(5):1050–1060. doi:10.1002/cpt.2488.
- **10.** Dowden H, Munro J. Trends in clinical success rates and therapeutic focus. *Nature Reviews Drug Discovery.* 2019;18(7):495–496. doi:10.1038/d41573-019-00074-z.
- **11.** European Medicines Agency. Reflection paper on forecasting demand for medicinal products in the EU/EEA. EMA/162549/2021. 2021.
- **12.** U.S. Food and Drug Administration. Drug Shortages: Root Causes and Potential Solutions. Report of the Drug Shortages Task Force. 2019; updated 2020.
- **13.** U.S. Food and Drug Administration. Drug Shortages Report to Congress: Calendar Year 2024. 2025.
- **14.** Fildes R, Goodwin P. Stability in the inefficient use of forecasting systems: a case study in a supply-chain company. *International Journal of Forecasting.* 2021;37(2):1031–1046. doi:10.1016/j.ijforecast.2020.11.004.
- **15.** Papathanasiou MM, Stamatis C, Lakelin M, et al. Autologous CAR T-cell therapies supply chain: challenges and opportunities? *Cancer Gene Therapy.* 2020;27:799–809. doi:10.1038/s41417-019-0157-z.
- **16.** Lee MK, Seyedmousavi S, Auvity S, et al. Forecasting the potential impact of cell and gene therapies in France: projecting product launches and patients treated. *Frontiers in Medicine.* 2024;11:1324602. doi:10.3389/fmed.2024.1324602.
- **17.** Truebel H, Seidler M. Mitigating bias in pharmaceutical R&D decision-making. *Nature Reviews Drug Discovery.* 2022;21(12):874–875. doi:10.1038/d41573-022-00157-4.