2026

# Biotech AI Report

## Breakthroughs, bottlenecks, and the power shift shaping biotech's AI future

Benchling

# Key findings

## 1 AI in biotech has found its first 'killer apps.'

A handful of tools have broken out of pilot mode and become everyday parts of R&D for biotech AI leaders. These are the use cases scientists now trust and rely on: literature review and knowledge extraction (76% adoption), protein structure prediction (71%), scientific reporting (66%), and target identification (58%). These early 'killer apps' succeed because they operate where data is clean and local, results are easily verifiable, and they fit naturally into a scientist's workflow.

## 2 AI is transforming R&D long before the first AI-designed drug.

The effects of AI are showing up earlier in the pipeline, where decisions about targets, constructs, and experiments set the trajectory for everything that follows. Half of those adopting AI in biotech already report faster time-to-target, 56% expect meaningful cost reductions within two years, and 42% see an uplift in accuracy and hit rates with scientific models. In a field where development takes 10-12 years, upstream improvements compound. Faster cycles, smarter decisions, and fewer dead ends matter enormously.

## 3 Biotech is hitting a ceiling with AI in complex, regulated science.

AI adoption drops sharply in areas like generative design (42% adoption), biomarker analysis (40%), and ADME prediction (29%) and IND submissions (24%). The limitation is more often the data environment than the models themselves. In these domains, data lives across a dozen systems, key metadata is often missing, and verifying outputs can take longer than the experiments themselves. Yet these are stages where decisions are complex and consequential, and where teams say they want AI to help next. Over the next two years, organizations expect to move from task-level copilots to systems that coordinate experiments and decisions end-to-end. The biggest areas of planned growth include AI for workflow orchestration, manufacturing optimization, multimodal models, and early co-scientist systems.

## 4 The #1 reason AI pilots fail is poor data quality and availability, and leaders are fixing it.

AI can design proteins, simulate processes, and reason across different types of data. But most organizations are running AI on systems built before this was possible. The top two reasons cited for pilot failure: data quality and availability followed by IP and compliance friction. The organizations using AI at higher rates look different; they've invested in connecting their data and integrating wet lab and computational workflows. This is the foundation AI needs to scale beyond its first wins, and get to the deeper, complex AI builds connecting cross-team data, workflows, and permissions.

## 5 Scientists have shifted with AI; the infrastructure needs to catch up.

AI is becoming a default tool: 89% use copilots or reasoning tools as their first stop for querying data. And 66% of respondents report an increase in trust in LLM outputs over the past year. Use of external data is accelerating — 77% rely on it, 71% increased usage this year — and scientists expect open-source tools and modern, flexible access.

## 6 Biotech creates its own builder culture to make AI work.

Model development, validation, and quality systems cannot be separated in a complex, regulated industry. This has led to a more focused builder culture: balancing internal development with external support, prioritizing speed, and relying on a new class of talent that can translate between science, AI, and business value. Two organizational priorities are emerging:

- **Scientific translators are on the rise.**
  Biotechs need people who can connect the dots between biology and AI, so they're developing this role in-house. Internal upskilling is the most common source of AI talent (67% citing), far outpacing hiring from tech (21%). Centralized AI groups with embedded R&D specialists are becoming standard, and AI leadership most commonly sits inside R&D.

- **Organizations are reorganizing for speed.**
  Leading biotech AI teams run AI accelerators, interdisciplinary sprint groups that test, validate, and fail fast. And a "build what differentiates, buy what scales" mindset is taking hold: teams build bespoke models where their biology is unique (55% citing), and more commonly buy proven components (60%). Only a third outsource development.

### We're seeing the beginning of what biotech's AI operating system could look like.

The foundations are forming: data flowing securely across teams, models guiding scientific decisions, workflows linking digital and physical experiments, and teams operating with both agility and scientific rigor. The organizations that invest in connected data, hybrid talent, and modern workflows will turn today's AI momentum into a continuous, compounding system of discovery.

# Data sources & methodology

This report draws on a November 2025 survey of ~100 biotechnology and pharmaceutical organizations actively using AI across research and development (R&D).

Importantly, this is not a general industry sentiment study. It is a view into the specific practices and priorities of biotech's AI leaders, the organizations that are already using AI regularly and shaping modern R&D.

All respondents are based in the U.S. and Europe, and represent a mix of scientists, technologists, and executives working in the following functions: discovery research, process and analytical development, bioanalytical science, and animal safety and toxicology. To qualify, respondents were required to be using AI in their organizations for R&D purposes today. Among those surveyed, 62% report using AI regularly across R&D, while 38% use AI in more targeted, highly specific use cases. The survey was conducted by an independent research firm and expert network to ensure objectivity and industry relevance.

**Demographics**

Industry

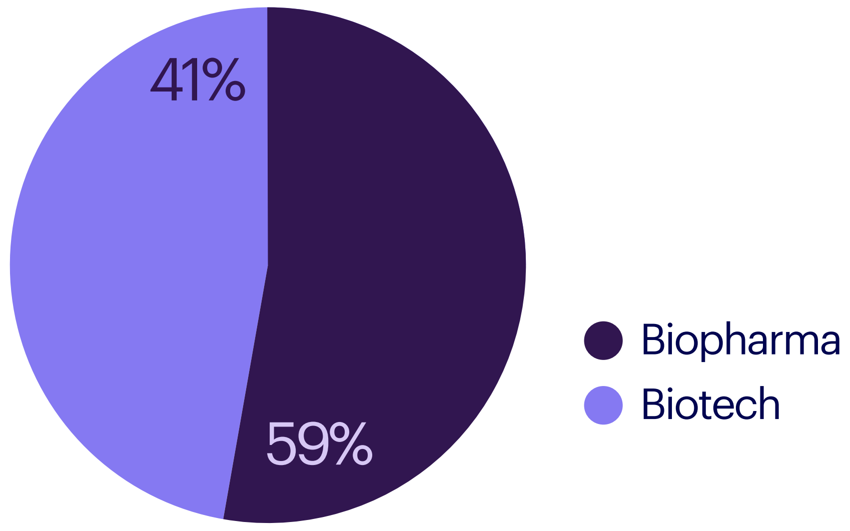


Region

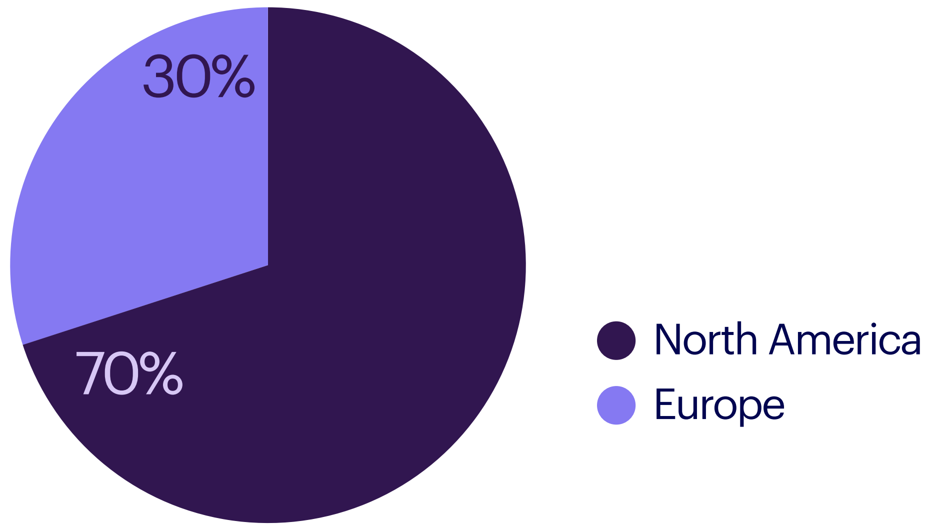


Company size

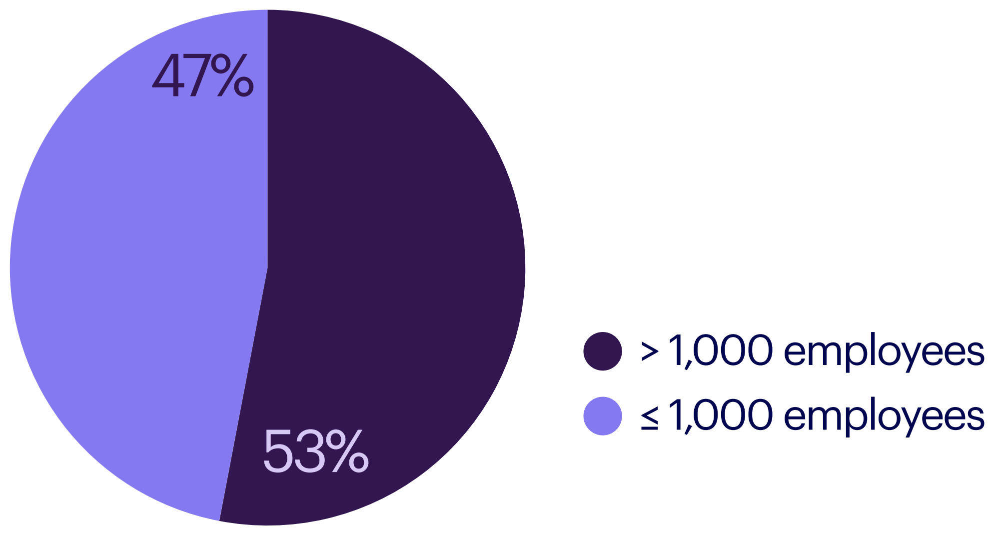


Role

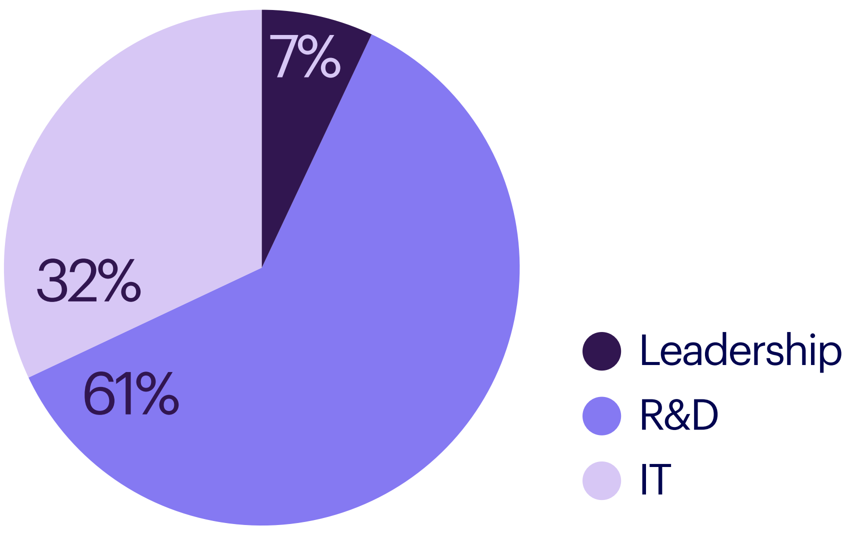

# AI in R&D

## AI is delivering real gains, but depth is still limited

Biotech's AI leaders have moved quickly: copilots and models are now improving speed, decisions, and hit rates. The industry's first 'killer apps' are here, the AI tools that scientists now use every day because they're reliable, accurate, and easy to validate.

But today's progress is mostly broad, not deep. Most use cases sit in prediction and clean, local data. The next wave — workflow automation, multimodal models, and deeper use in development — demands more integration and stronger data foundations.

# Biotech AI leaders show broad adoption across copilots, agents, and models; depth comes next

Biotech AI leaders are using copilots and assistants widely across day-to-day R&D operations, from documentation and reporting to search and analysis. Scientific model use is strong but more distributed, with 27% still piloting. AI use remains higher in non-scientific workflows (59% use regularly) than in scientific ones (44%), showing AI use cases are concentrated in simpler, low-friction, predictable workflows.

### Adoption for AI applications and scientific models R&D

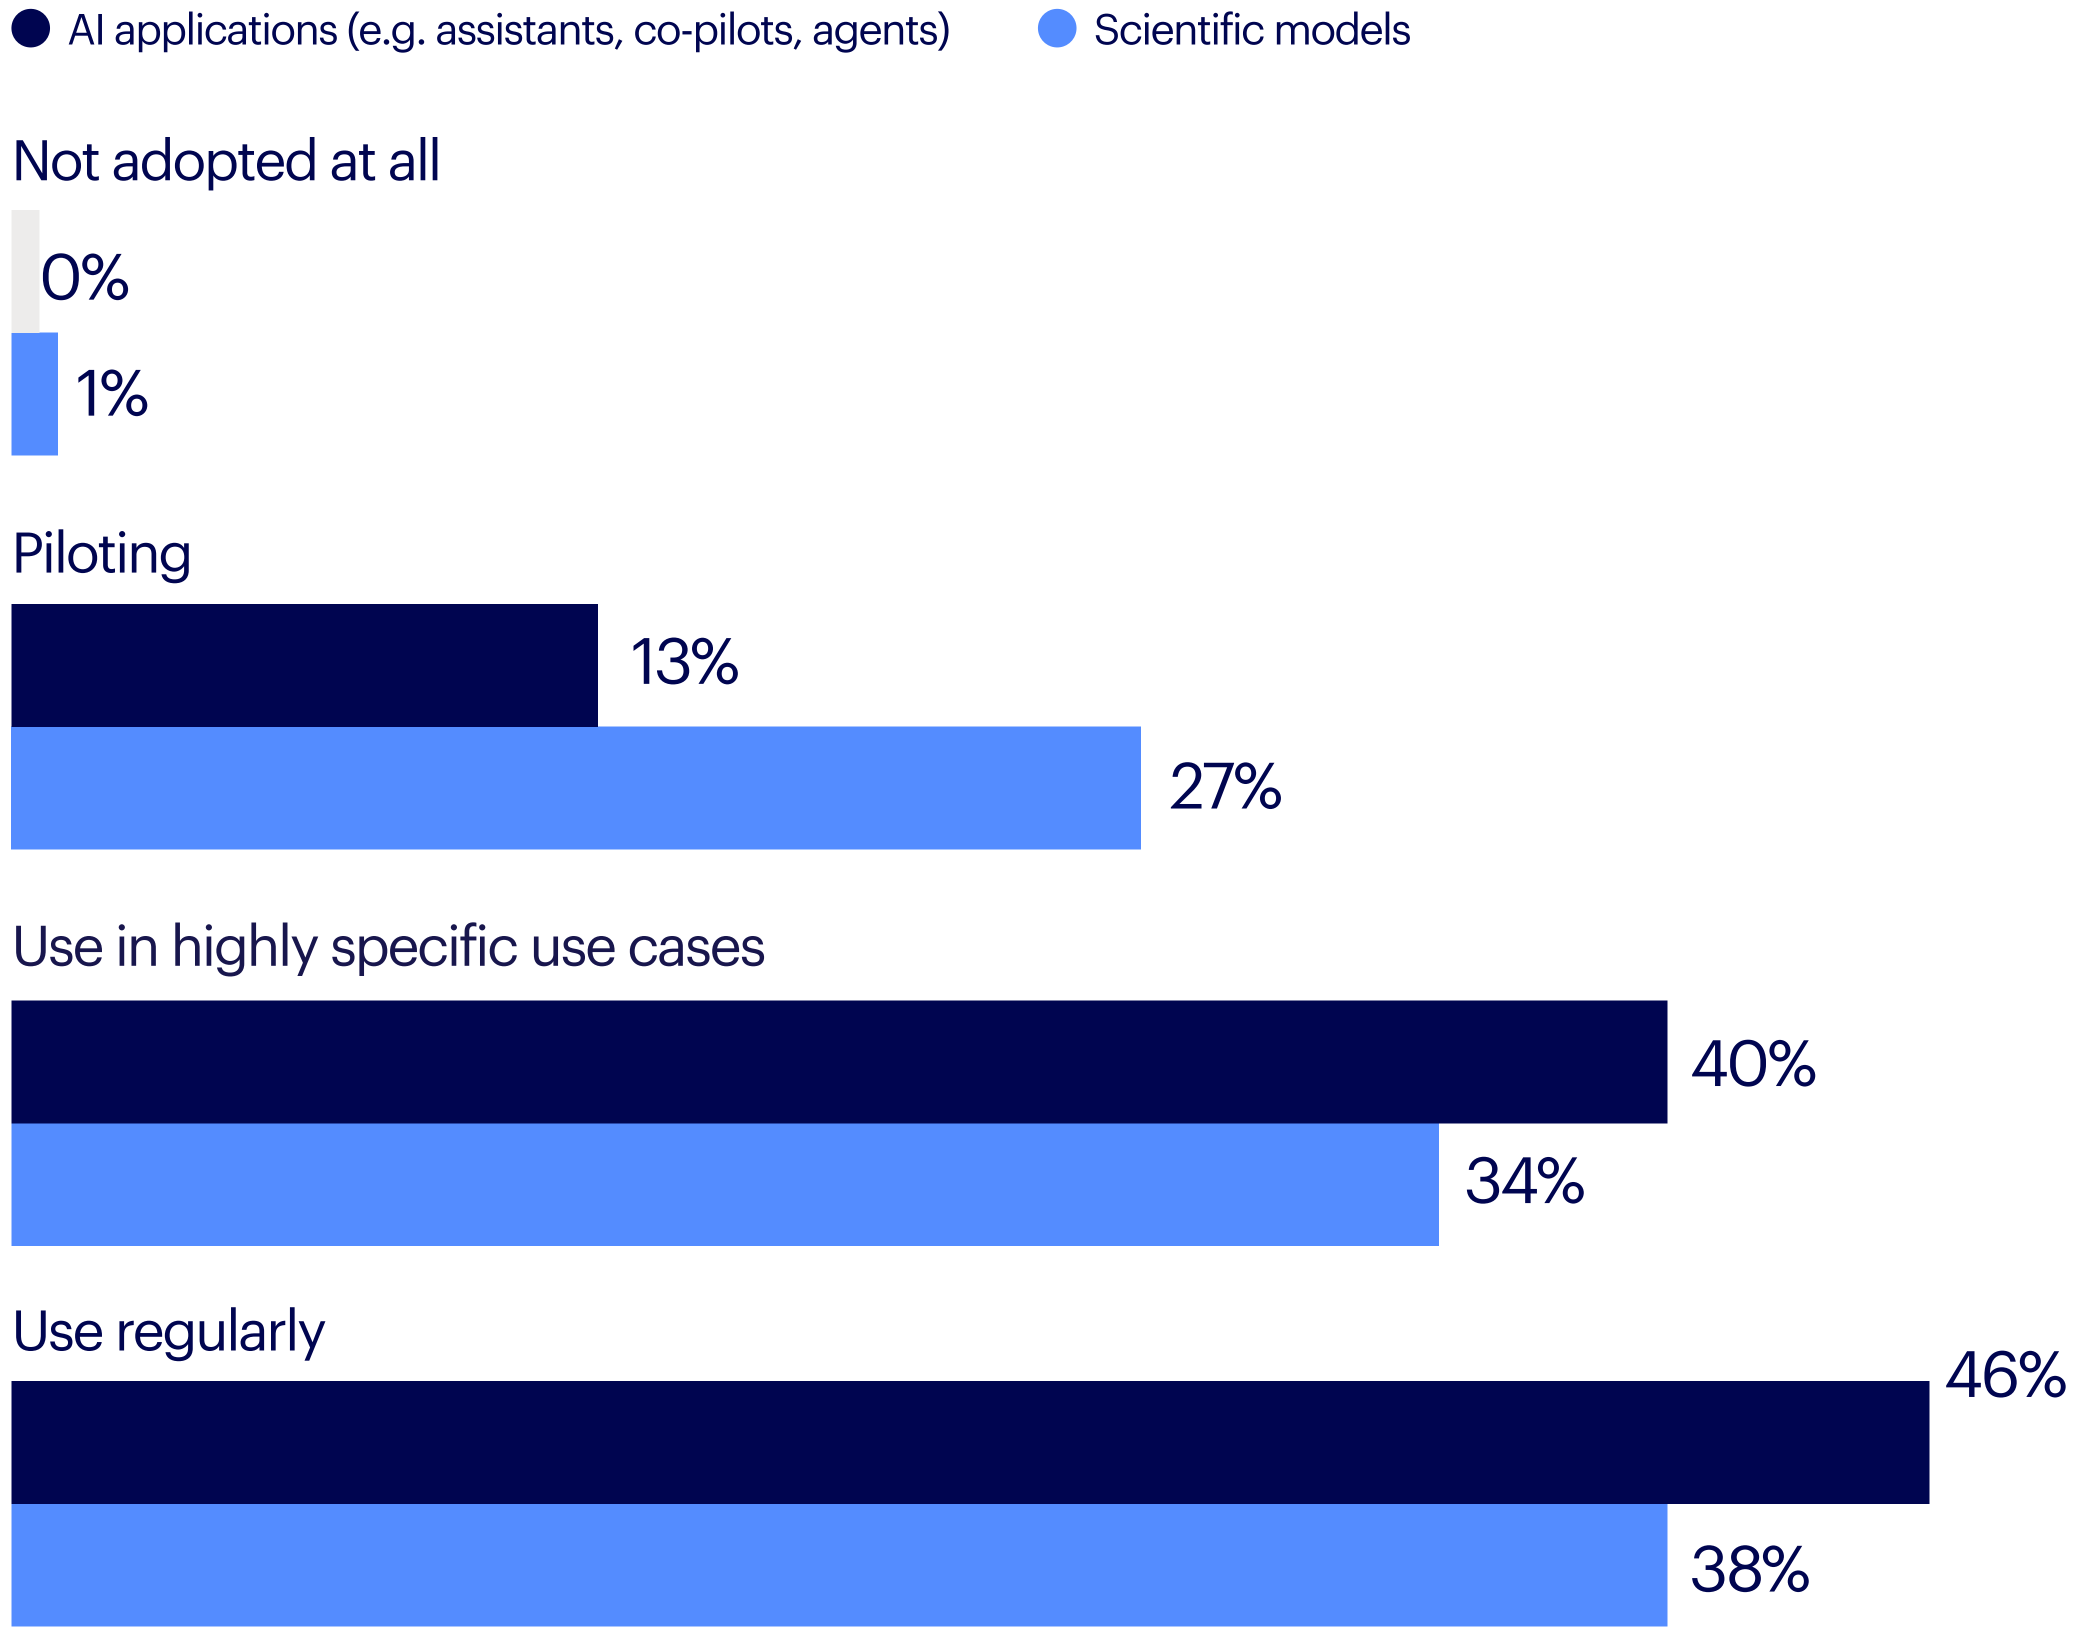


Results from all respondents (N = 104).
Numbers do not add to 100 due to rounding

### AI adoption for scientific vs non-scientific use cases

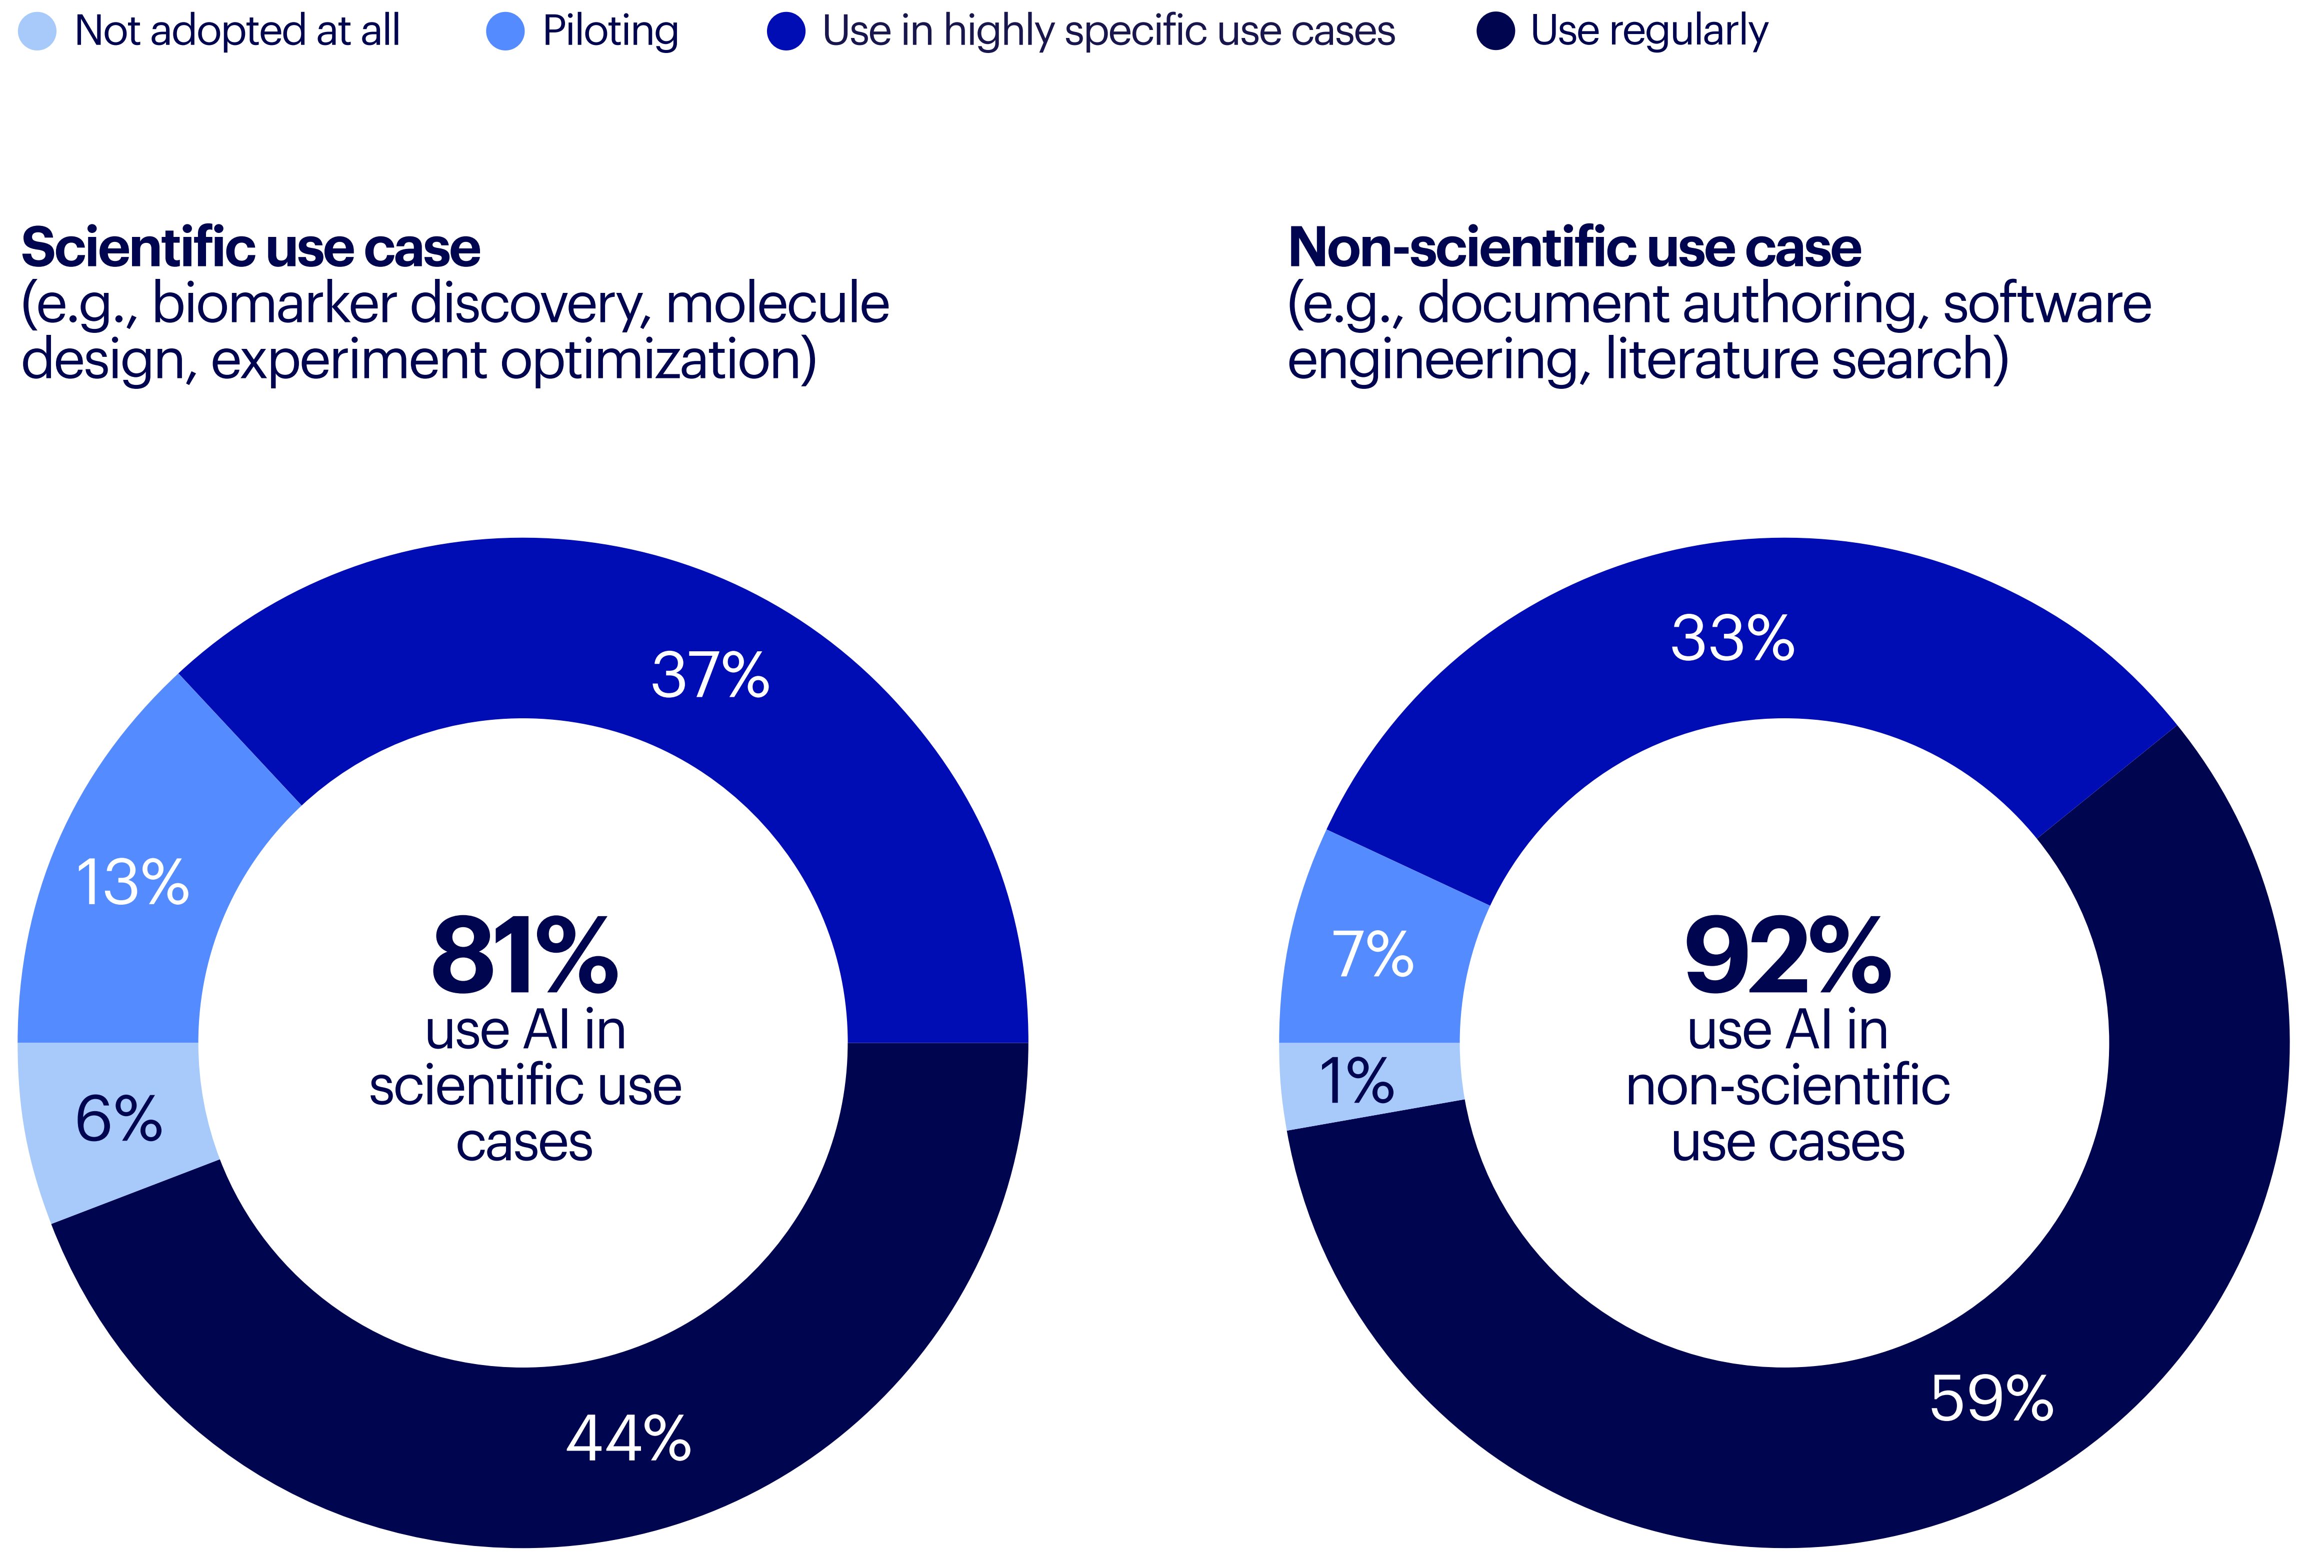


Results from all respondents (N = 104).

AI in research, discovery

# Biotech's first 'killer apps' take hold where data is local, clean, and easily validated

AI adoption is strongest in areas with structured or text-based data that's easy to validate. The first killer apps include literature and knowledge extraction (76% adoption), protein structure and property models (71%), scientific reporting (66%), and target identification (58%). Adoption drops as the science gets more complex. Use cases like generative design (42% adoption), biomarker analysis (40%), hit ID (37%), and ADME (29%) rely on multimodal, cross-team data and tight wet-lab feedback loops, conditions that slow progress today.

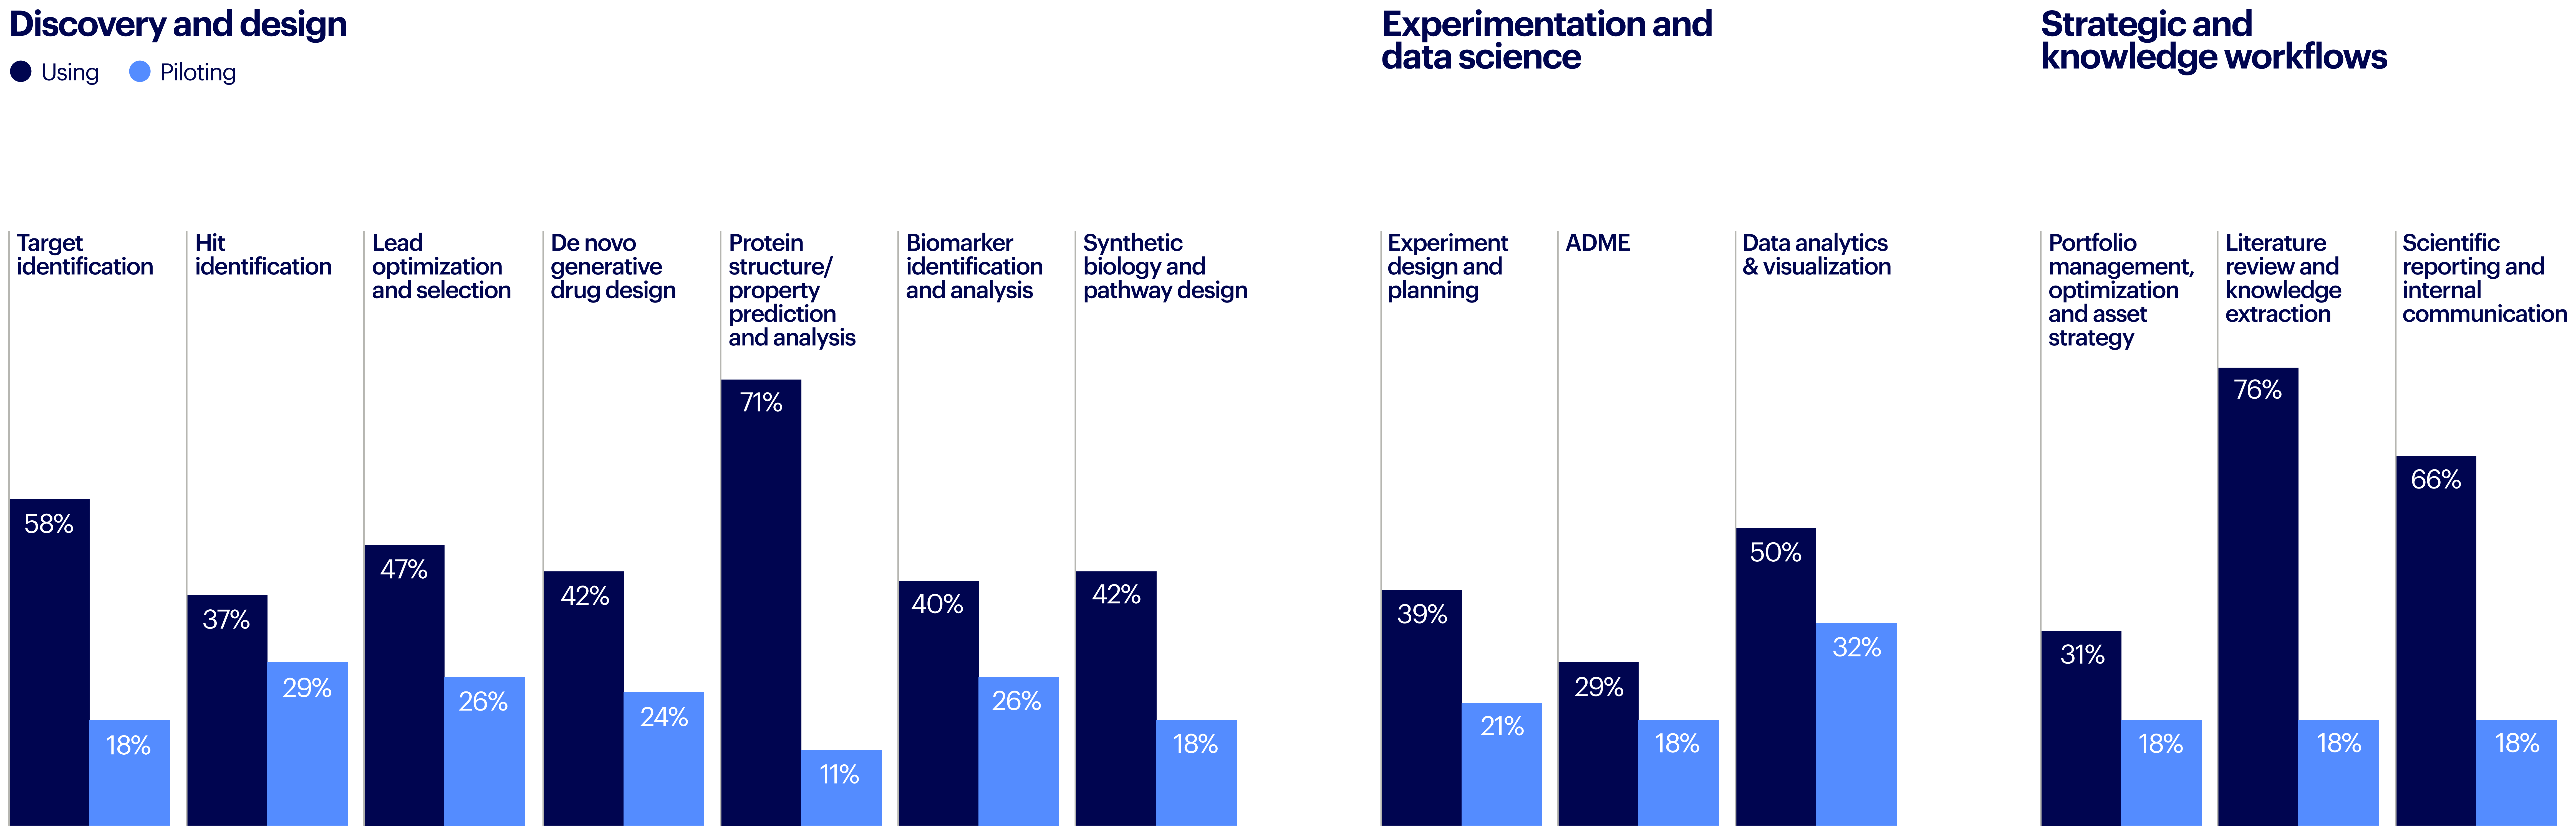


% indicates adoption of AI for each category, N = 38

# AI in drug development lags, early progress is in knowledge and documentation work

Early wins cluster in documentation-heavy, process-driven tasks such as knowledge management (58% adoption), data and audit log review (48%), and tech transfer (48%), places where copilots and retrieval models can operate on clean, compliant data. AI adoption in downstream development still lags research and discovery. The hurdles are operational: stricter compliance requirements, limited access to AI or data-science support, and development data that spans many teams and systems, making it harder to connect and automate.

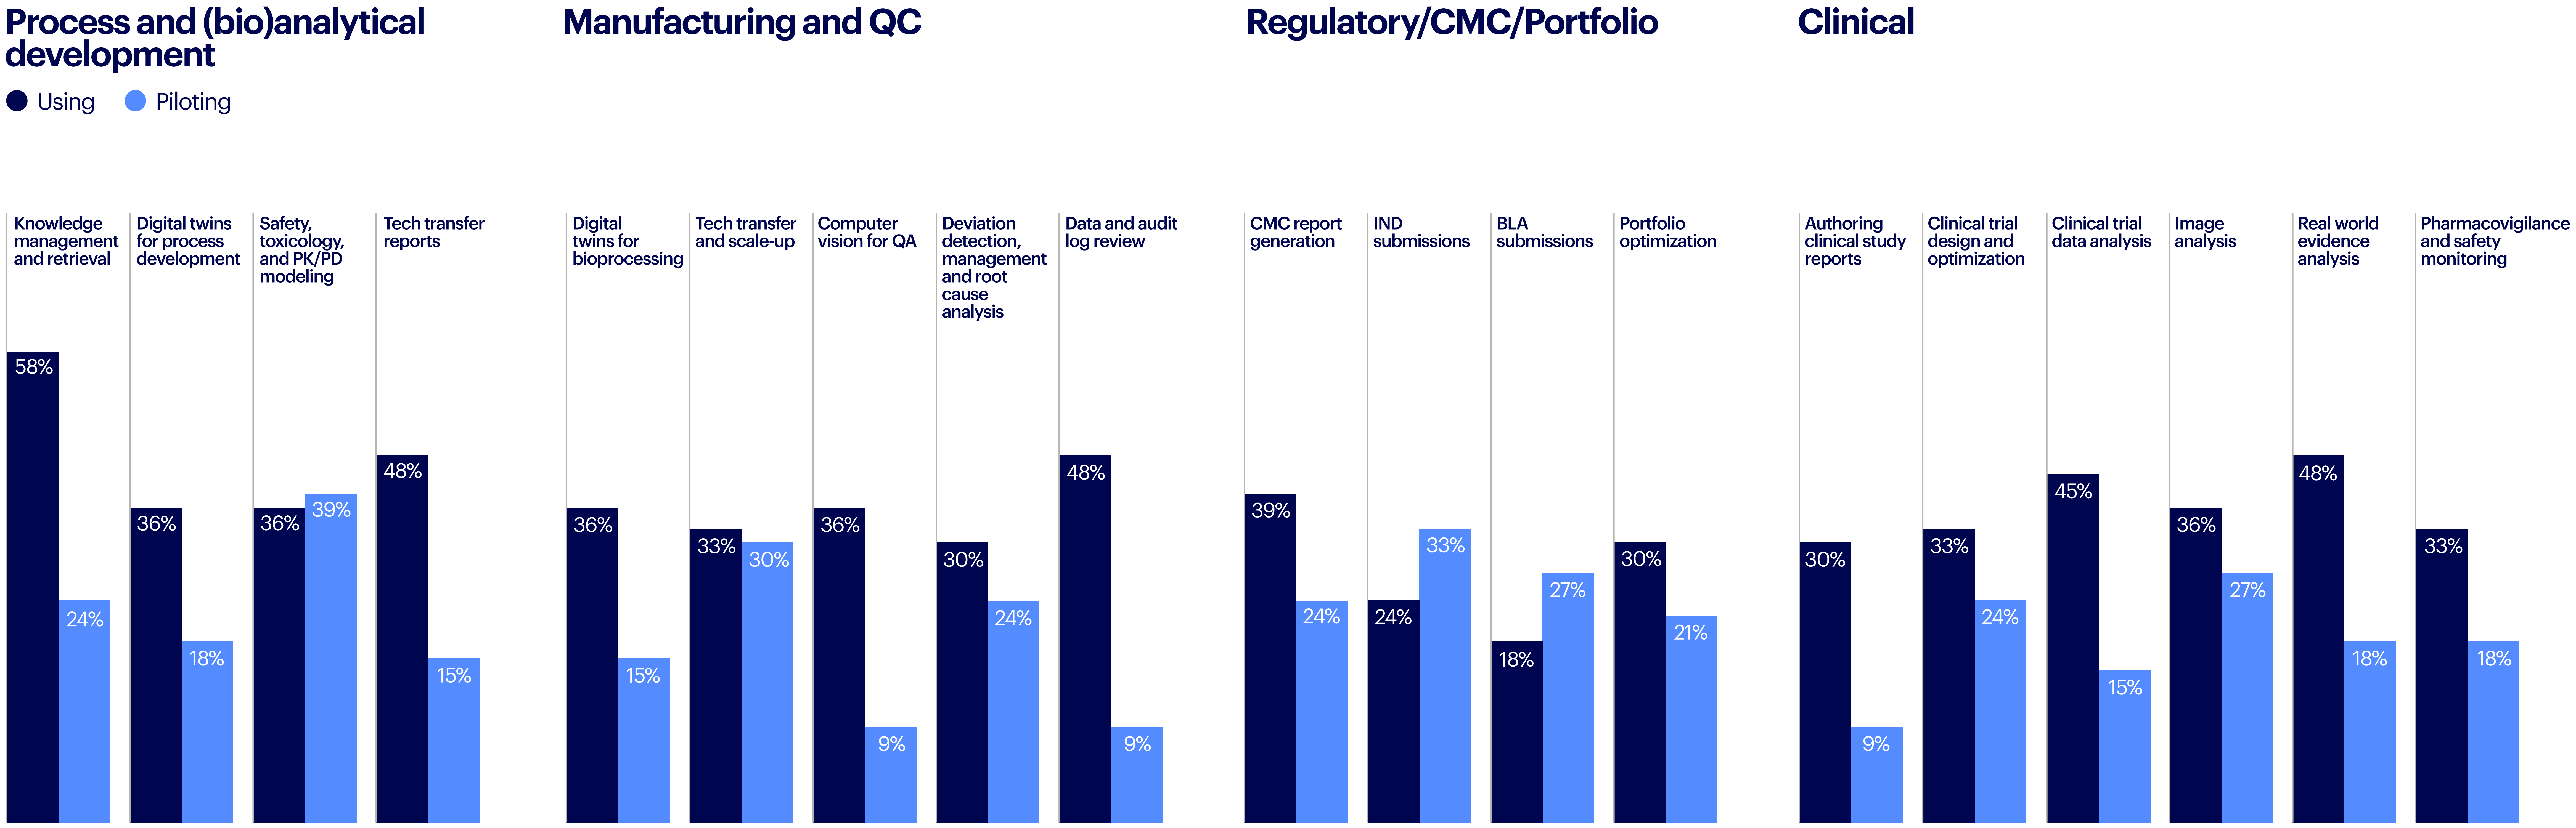


% indicates adoption of AI for each category, N = 33

Scientific models

# Prediction dominates today; multimodal and generative models are the next frontier

Predictive models lead adoption because they sit on mature, well-structured datasets and have fast validation loops: 73% use structure prediction and 52% use docking models. These avoid common pitfalls like missing metadata, slow data generation, and fragmented assay context.

Generative design is gaining traction but remains early (36% adoption), reflecting the need for richer, cross-team data to validate new sequences and designs. Specialized biological models — RNA, genome, single-cell (~22%) and multimodal Bio-LLMs (16%) — are rising but still limited by data availability and context.

Scientists also rely heavily on general-purpose LLMs (72%), reinforcing that usability, accessibility, and broad reasoning often outpace domain-specific sophistication in day-to-day work.

## Scientific model adoption

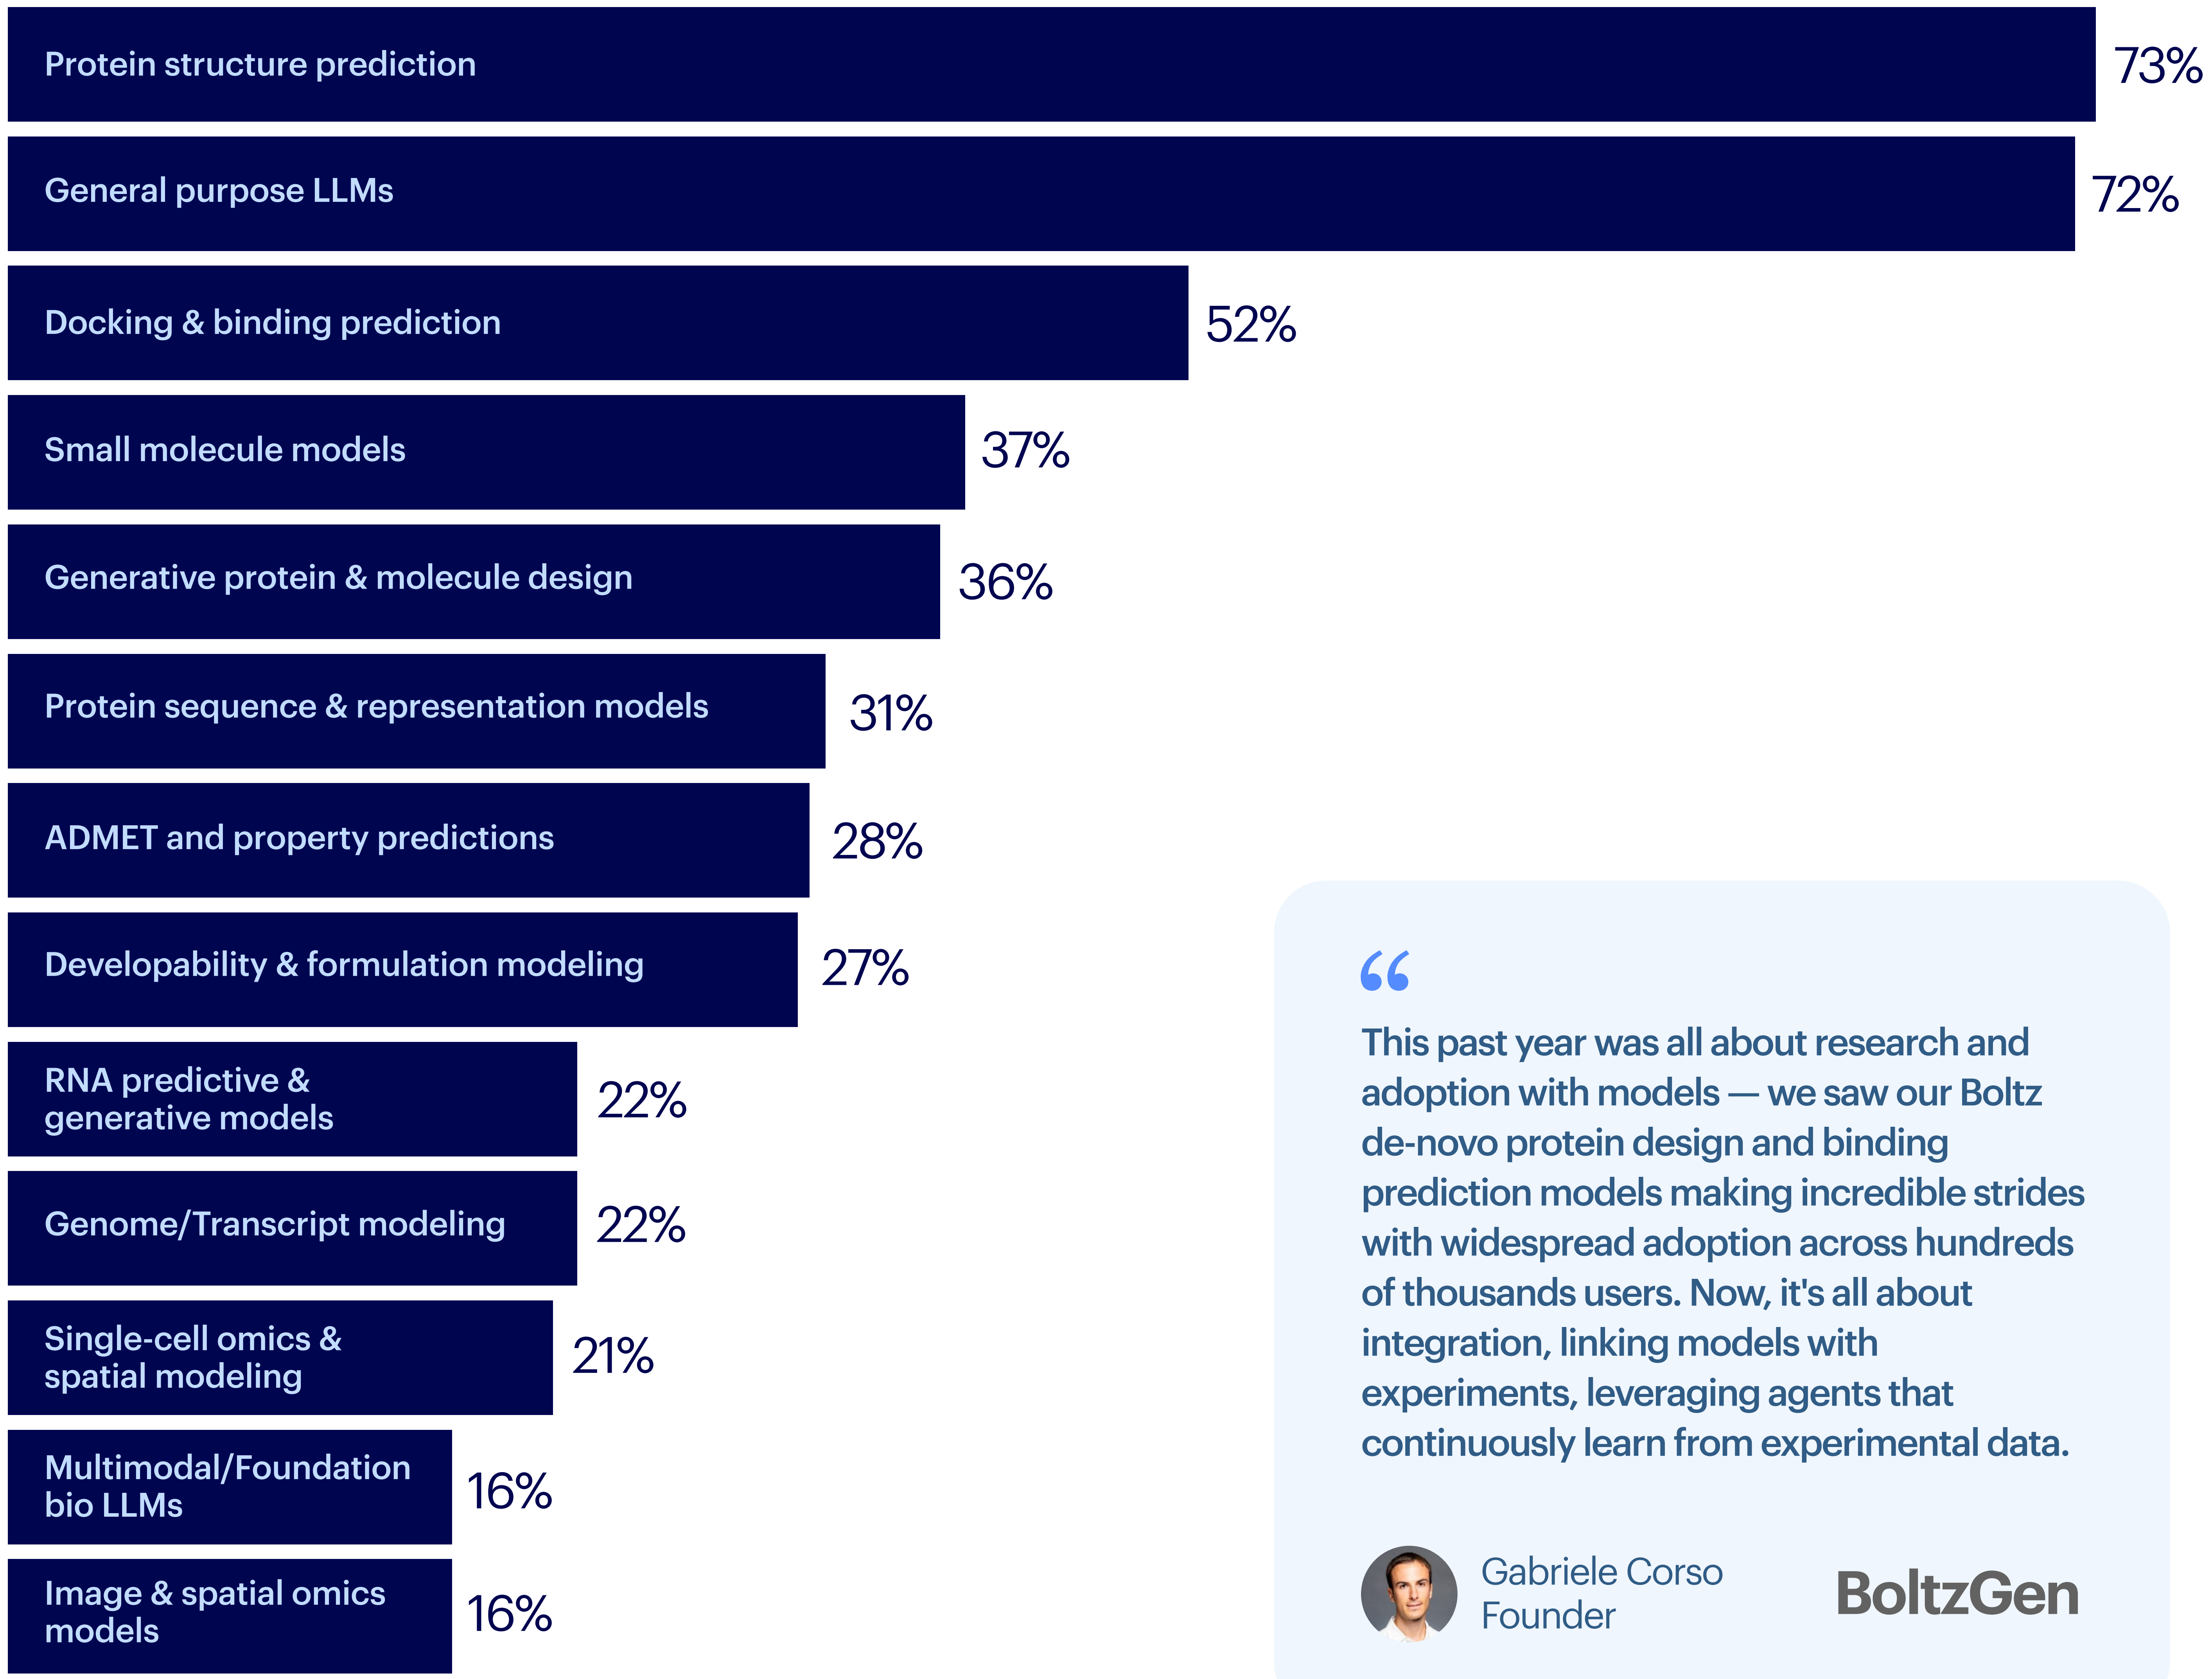


% indicates adoption for each category, N = 99

> "This past year was all about research and adoption with models — we saw our Boltz de-novo protein design and binding prediction models making incredible strides with widespread adoption across hundreds of thousands users. Now, it's all about integration, linking models with experiments, leveraging agents that continuously learn from experimental data.

Gabriele Corso
Founder

BoltzGen

# AI skews task-specific, not yet end-to-end

“

We’re seeing AI make its biggest impact in pharma when technology, science, and process design work in unison. At BMS, AI is already supporting nearly every facet of our work. Our integrated approach connects data, AI/ML, wet lab, and clinical expertise into one ecosystem, where insights continually inform decisions, accelerate learning, and help us discover and develop new medicines.

Matteo diTommaso
SVP IT R&D

Bristol Myers Squibb

“

There's no question that AI has raised the floor in biotech almost overnight, making molecular design and analysis dramatically faster and cheaper. But efficiency alone won't lead to fundamentally different outcomes. As AI becomes table stakes, the real advantage will come from teams that power it with data from proprietary assays and tight experimental feedback loops, allowing models to learn directly from relevant biology.

Gleb Kuznetsov
CEO

ManifoldBio

Most organizations use several agents and model types today, but very few have stitched them together into true, end-to-end pipelines. A scientist might use one model to extract insights from papers, another to generate a hypothesis, and a third to design an experiment, with each step existing independently.

As multimodal models and open standards like MCP (model context protocol) mature, teams want these tools to talk to each other, forming a connected, interoperable AI ecosystem. The end goal is data, decisions, and experiments flowing automatically from one step to the next.

**39%** use AI for **single** or task-specific AI e.g., protocol summarization, document generation, assay planning

**22%** use AI spanning **multiple** R&D workflows across teams or functions

**38%** use **both** single-use and cross-functional AI applications are in use or exploration

N = 104
Numbers do not add to 100 due to rounding

# AI shows significant impact on R&D, long before the first AI drug

In a field defined by decade-long timelines and high failure rates, AI shouldn't be judged by the first "AI-designed molecule," but by the hundreds of decisions and workflows along the way. And on those fronts, AI is already delivering. With AI applications (including copilots, agents, LLMs), 50% of biotech report faster time-to-target, 37% see better scientific outcomes like improved hit rates, and 56% expect cost reductions within two years as automation and agentic workflows scale.

> "
> AI is industrializing science by codifying the complexity of biomolecular systems into predictive and generative models. We're moving from trial-and-error experimentation to AI-driven design. By tightly coupling AI design systems with the lab, researchers are creating a powerful data flywheel that are shrinking drug discovery timelines from years to months. Together, these advances are scaling medical discovery and accelerating breakthroughs across the industry.

Kimberly Powell
VP of Healthcare
NVIDIA

## Measurable impact of AI applications on R&D

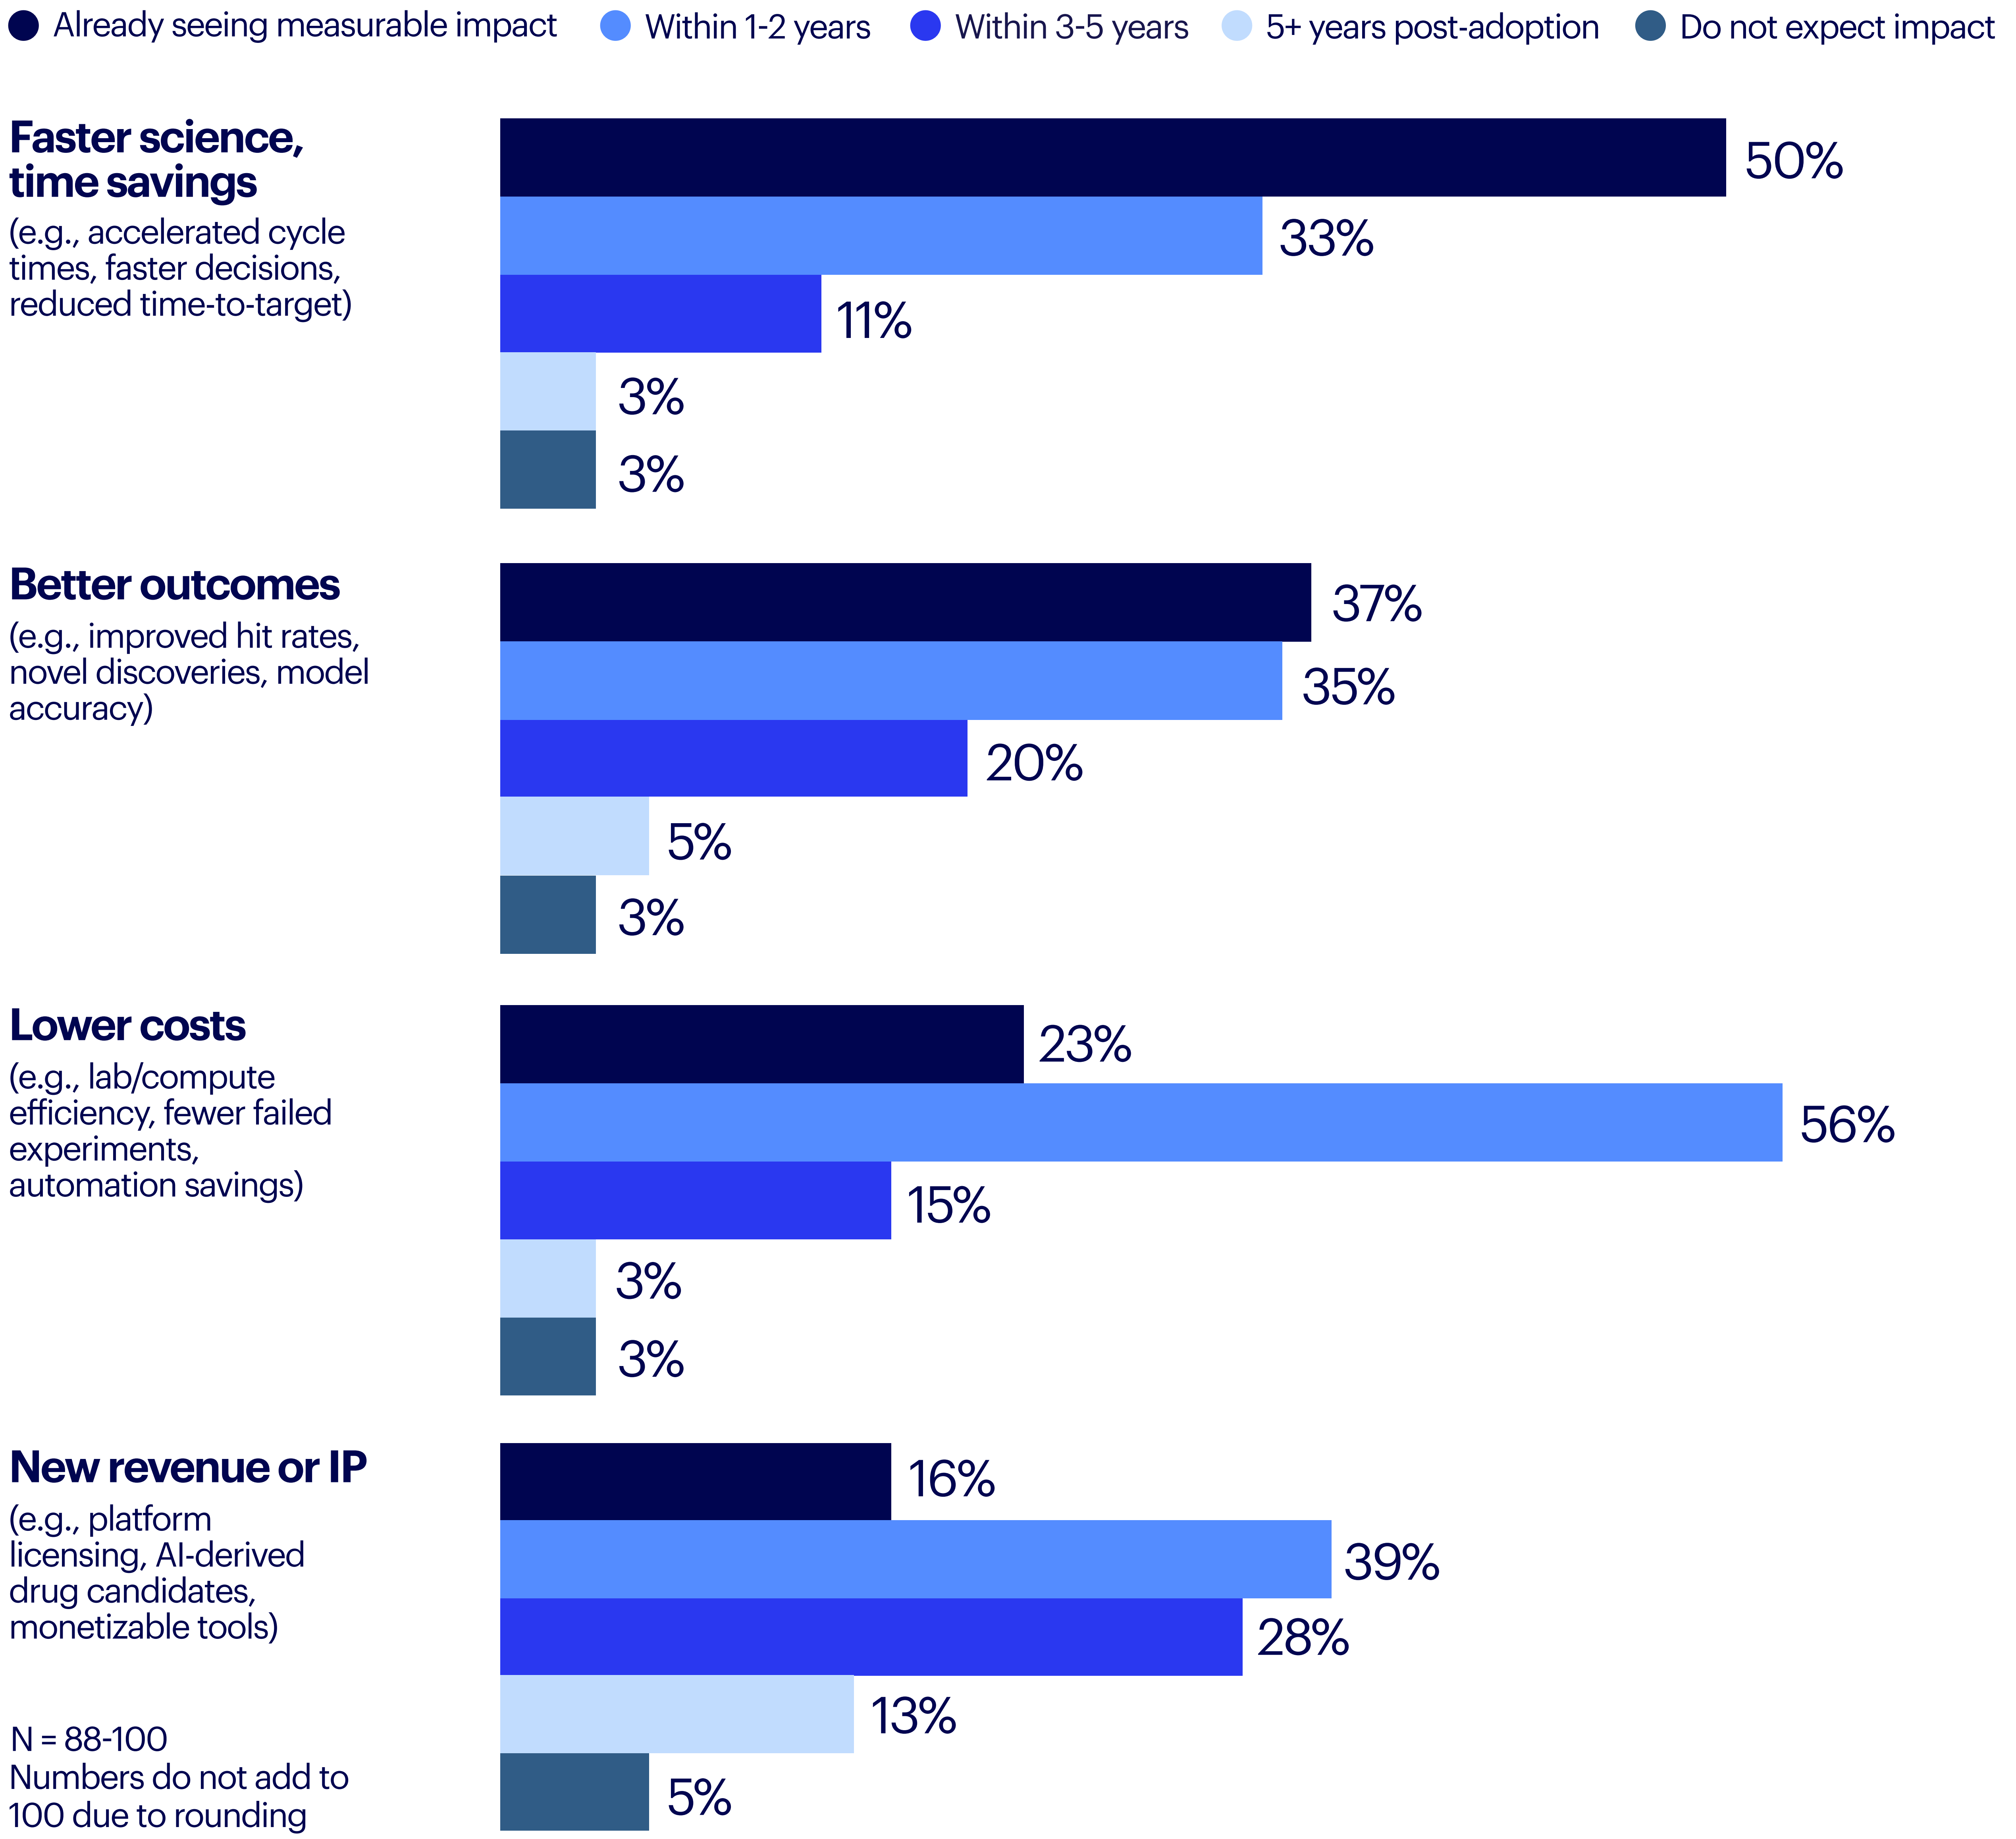


N = 88-100
Numbers do not add to 100 due to rounding

# Scientific models outperform on scientific quality and better outcomes

Biotech using scientific models are already seeing meaningful gains: 42% report improved hit rates or prediction accuracy, with another 42% expecting further improvements within 1–2 years. These benefits come alongside faster R&D (50% cite time savings today) and expected cost reductions (44% within two years). Revenue and IP impacts will take longer, with most projecting AI-derived assets on a 3–5 year horizon.

> "AI is fundamentally reshaping biotech R&D along multiple frontiers. As we see across other industries, biotech companies are deploying AI and AI agents to optimize business processes and workflows. More profoundly, we're witnessing the emergence of specialized AI tackling core biotech challenges like designing molecules, decoding complex biological systems and delivering medicines to the right patients in ways that weren't possible before. Together, we hope these developments will accelerate progress and improve the PTRS (probability of technical and regulatory success) of the biotech industry.
>
> Hetu Kamichetty
> Co-founder and CTO
> xaira

## Measurable impact of scientific models on R&D

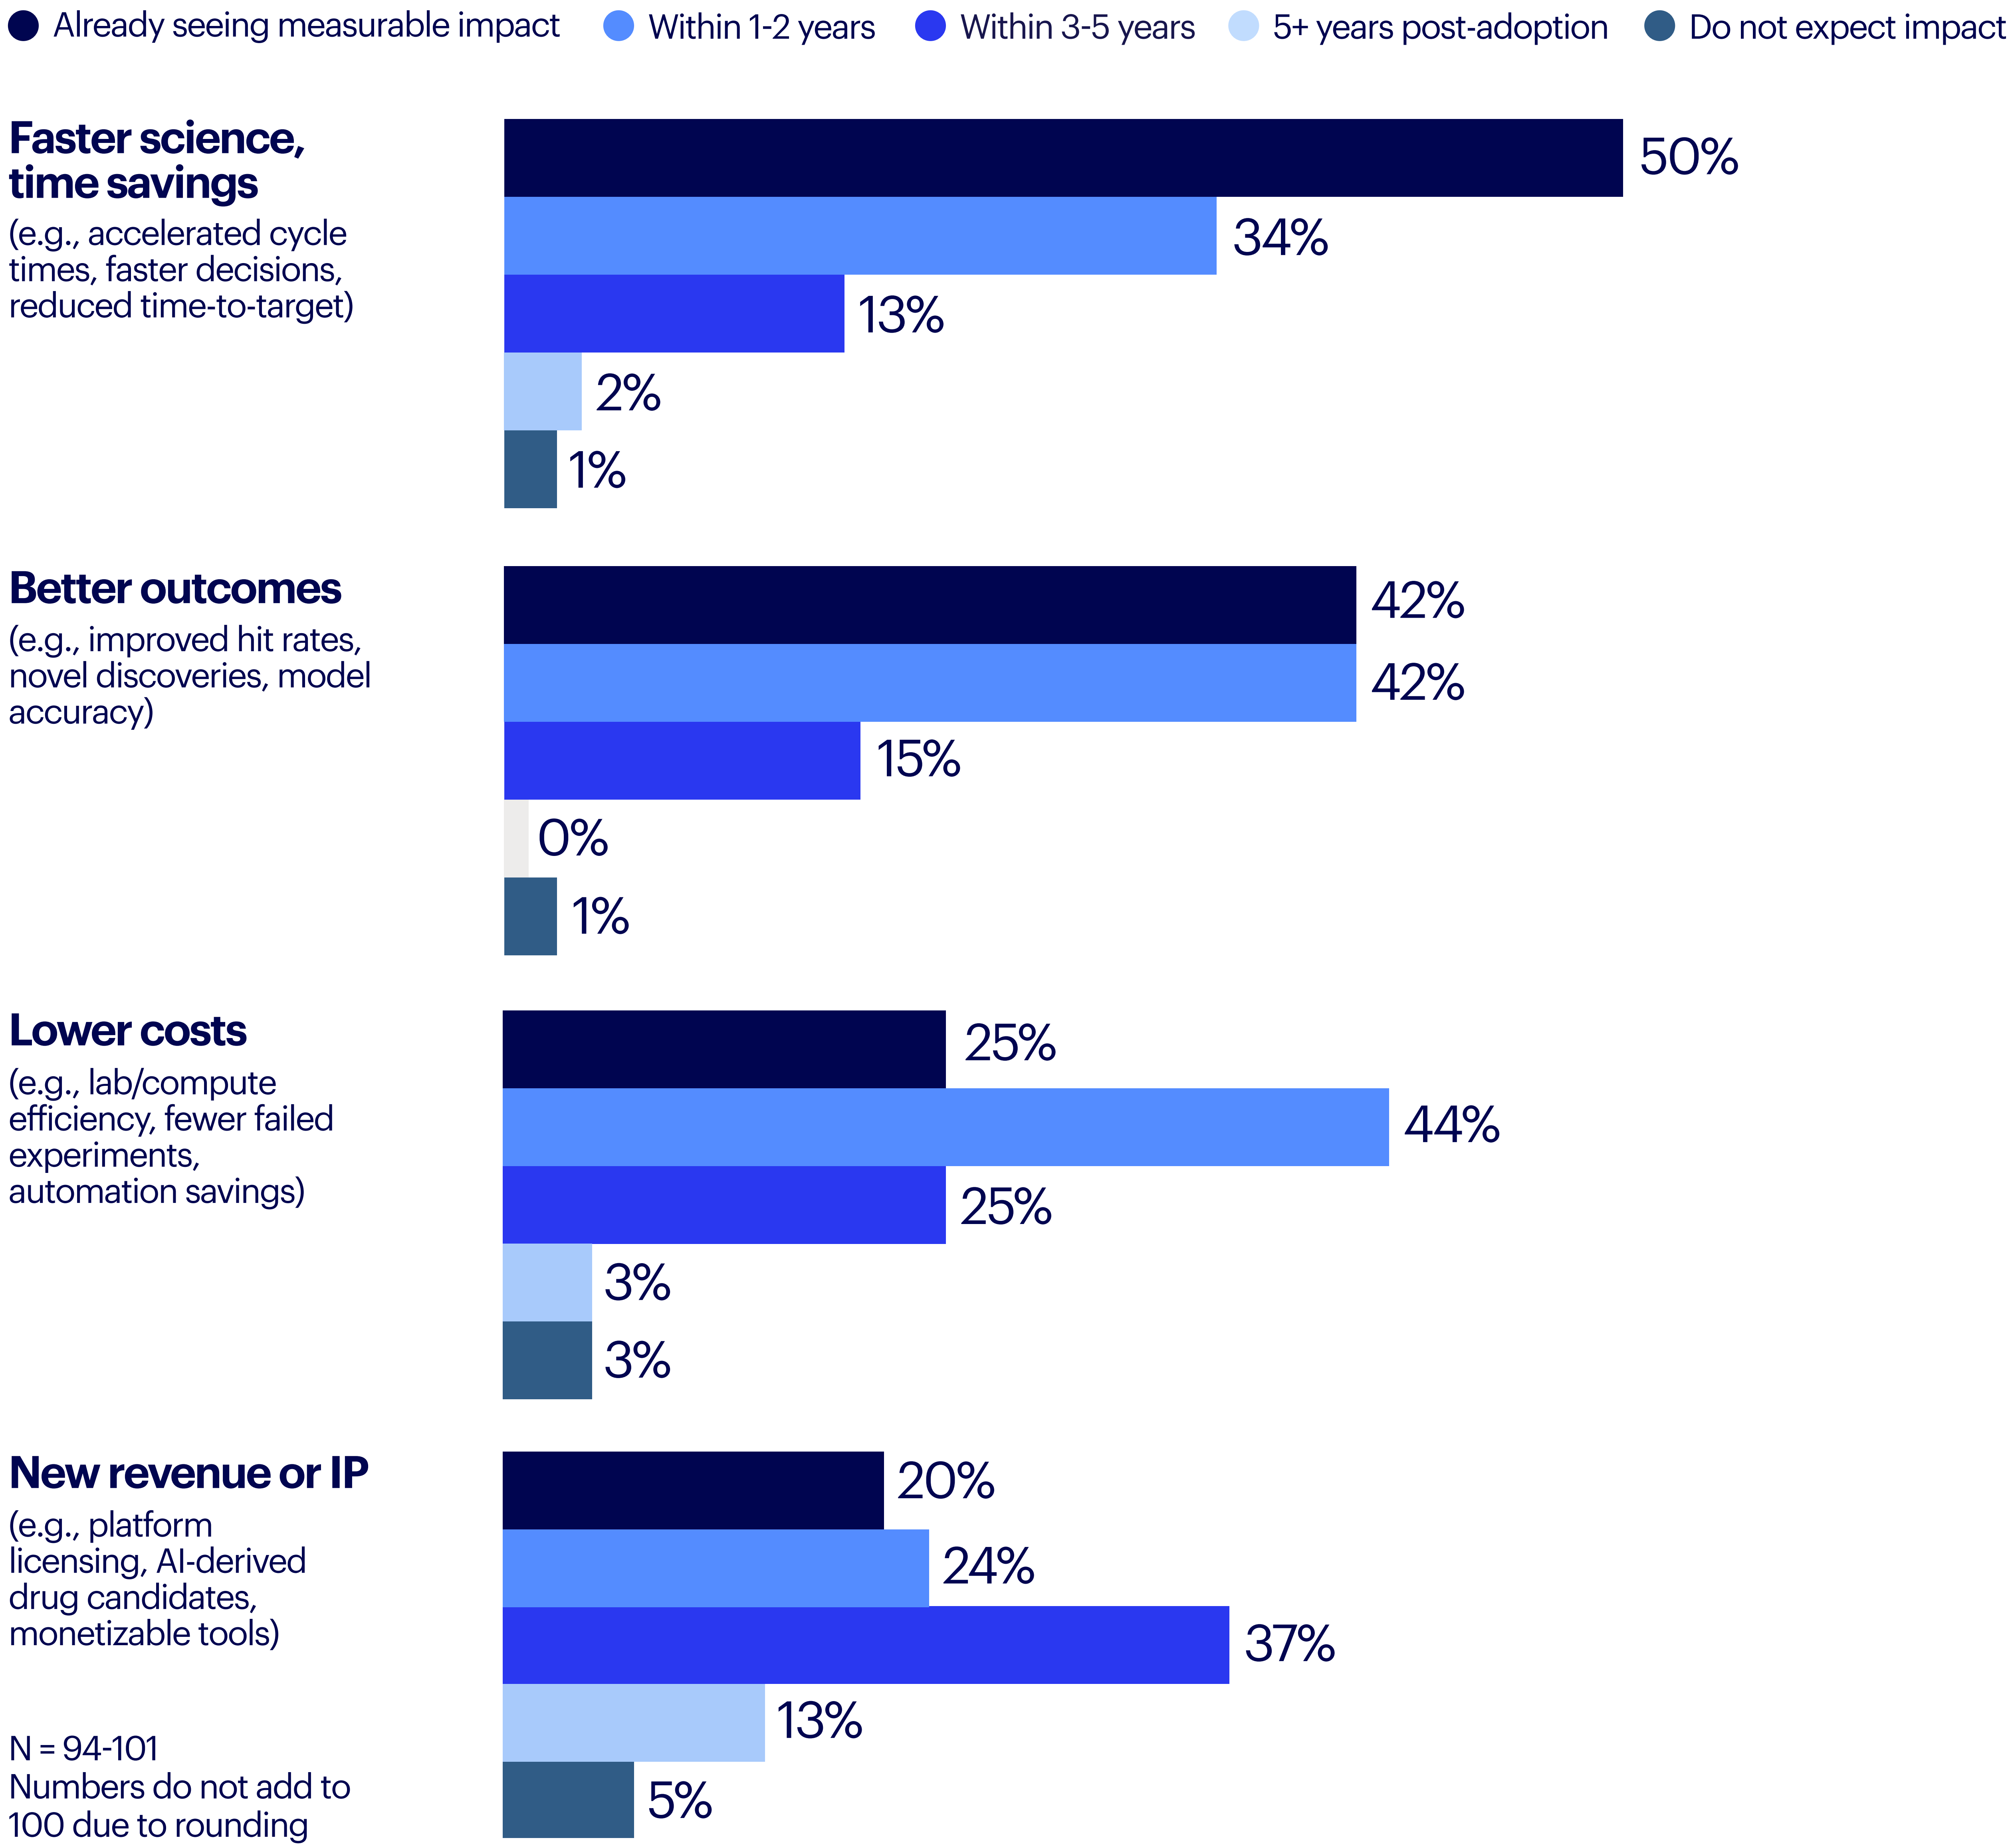


N = 94-101
Numbers do not add to 100 due to rounding

# AI's next wave: Moving toward autonomous, closed-loop R&D systems

Biotech's early AI wins came from copilots and task-specific tools. The next wave is more ambitious: AI that links teams, data, and instruments so experiments can run in tighter, semi-autonomous loops.

Over the next 1–2 years, biotech plan to adopt capabilities that require data to flow across teams and systems: workflow orchestration (38% plan to adopt in 1-2 years), manufacturing and supply chain AI (38%), multimodal models (37%), and AI for clinical development (33%). Only 26% use co-scientist agents today, another 32% plan to adopt them in the next two years, a shift from task assistants to agent chains that coordinate experiments, automation, and decisions.

### Which AI capabilities organizations plan to adopt, and when

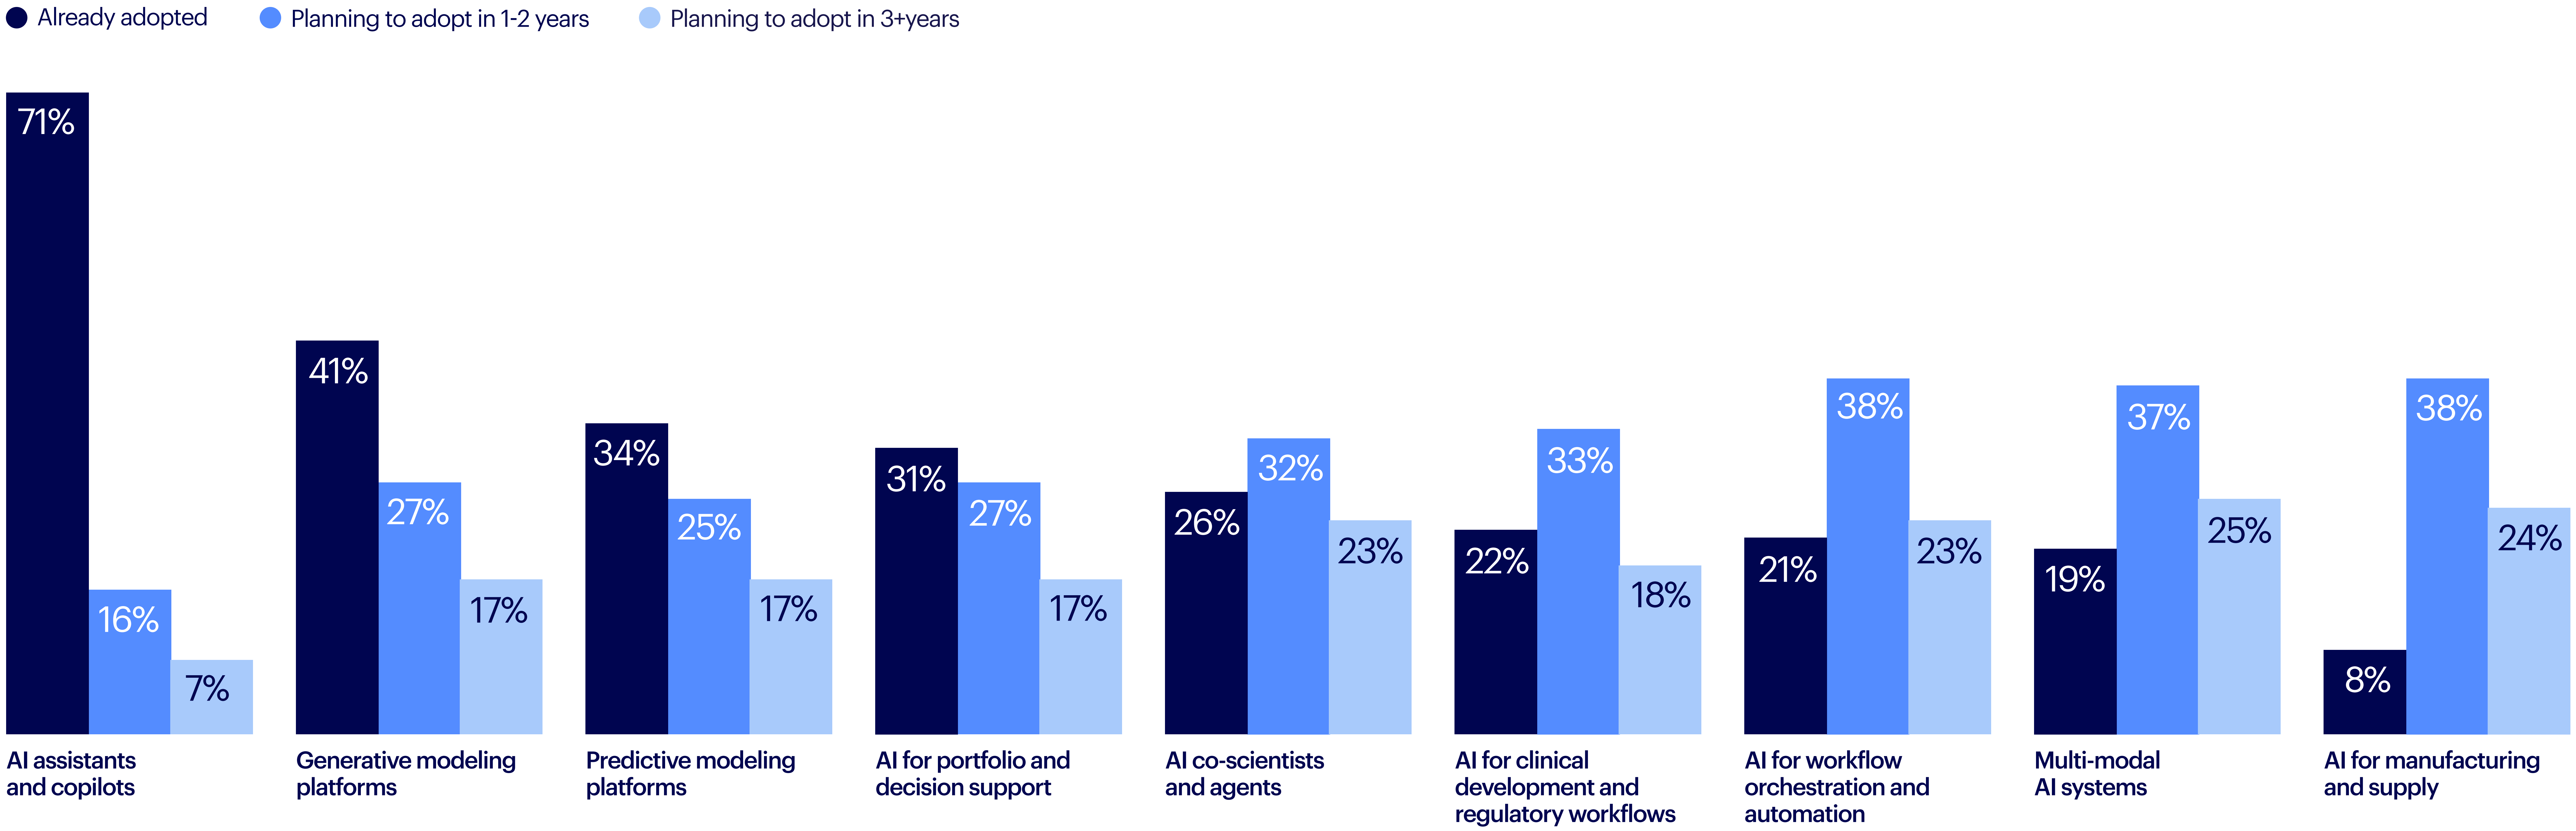


N = 89-102, but "adopted and abandoned" and "not sure" responses have been omitted.

# Pilot failures show AI is hitting the limits of data and biotech infrastructure

> 55% rate data challenges and 50% rank IP, security and compliance as major factors for AI pilot failure or underperformance. Lack of internal AI talent is not a significant factor.

Unstructured, multimodal biological data is biotech's biggest bottleneck and greatest resource. Biotech need ways to clean, structure, validate, and govern their data so downstream and cross-functional AI workloads actually work. The use cases are everywhere: connecting assay results to computational predictions, linking biomarker data across studies, or automating design-build-test-learn cycles. But none of it scales without the data and infrastructure to support it.

"

**Talking about 'AI for drug discovery' is as vague as saying 'chemistry for drug discovery.' What matters is whether we're generating the right measurements in the first place. Biology's data is often too messy or incomplete to teach machines what we want them to learn, and no amount of retroactive normalization can fix that. Progress comes when we invest in prospective data, the high-quality, well-annotated measurements that models can truly learn from.**

Lindsay Edwards
CTO

RELATION

## AI pilot failures or underperformance challenges

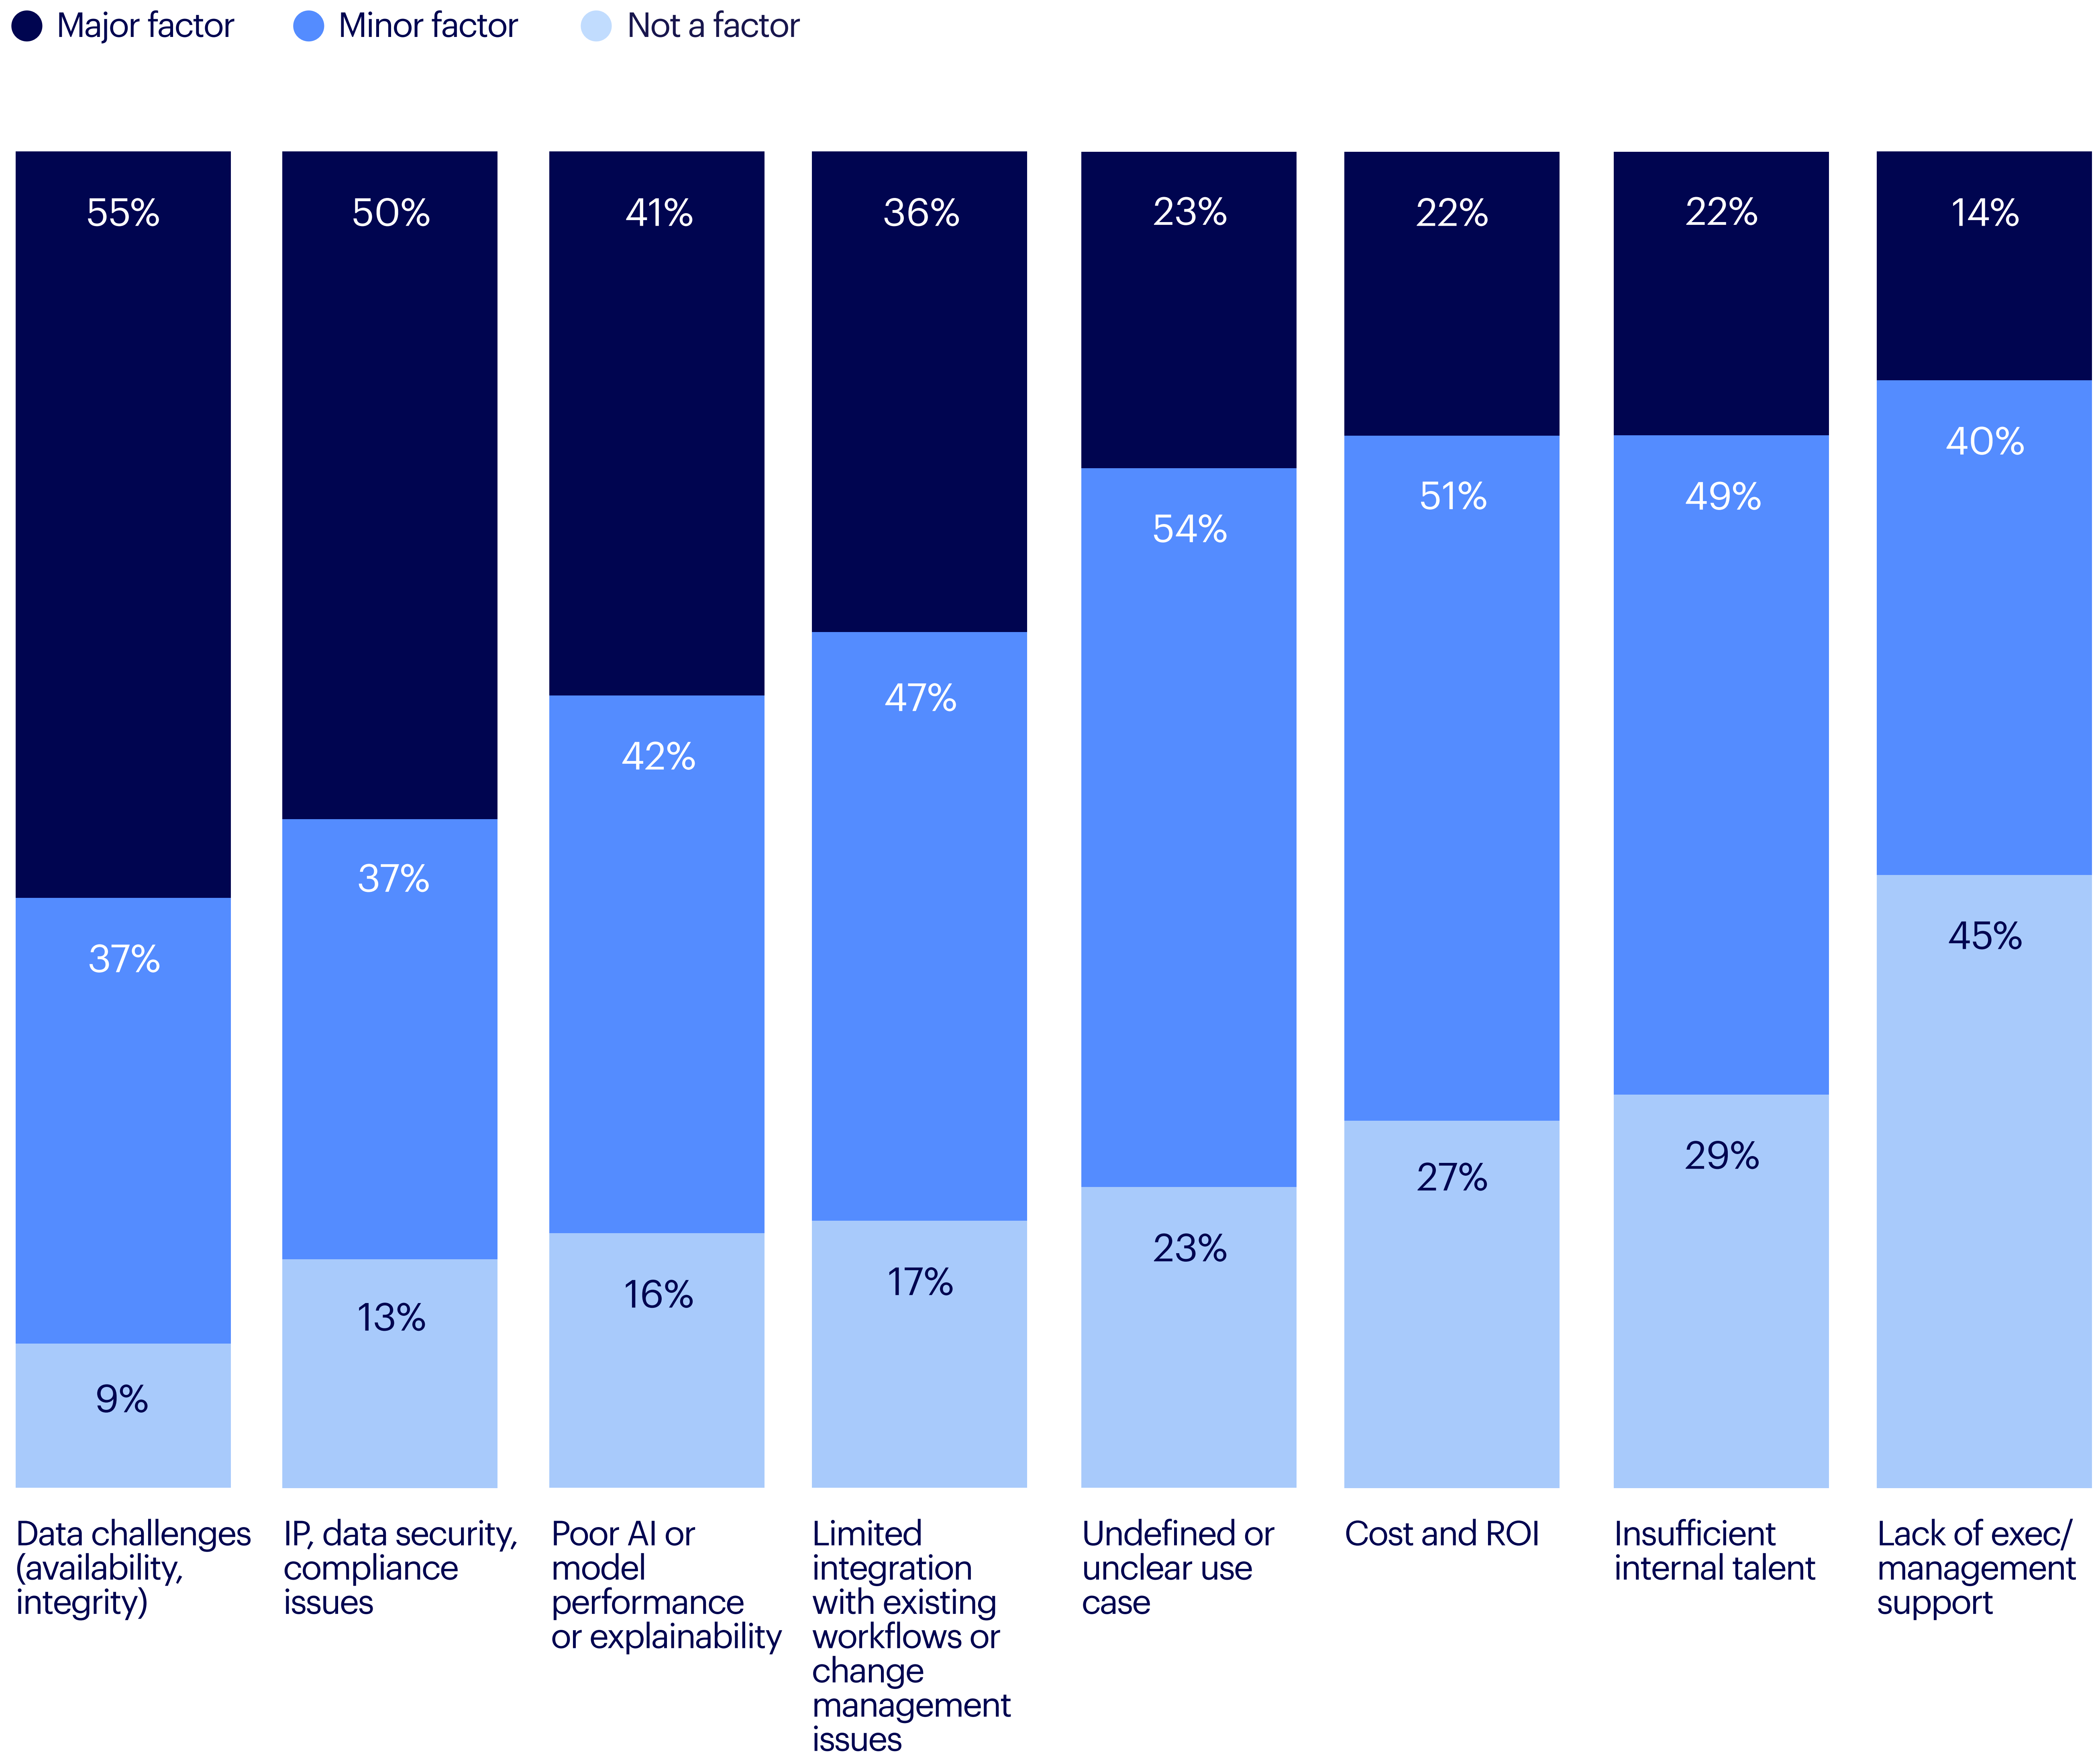


N = 104
Numbers do not add to 100 due to rounding

# Data & Infrastructure

## AI-first science needs an AI-ready infrastructure

AI is ready, but most R&D systems weren't built with AI or autonomous science in mind. Static, siloed data environments that were "good enough" a decade ago are now the biggest bottleneck. As scientists shift to AI-first workflows — with copilots as their primary interface, external and open-source data skyrocketing, and early forms of closed-loop, automated experimentation emerging — the leaders are building a more dynamic data foundation.

# External data use is accelerating, and it's reshaping what AI can do

AI becomes far more powerful when models can reason across broader, more varied data. Biotech are expanding beyond internal pipelines and incorporating public, commercial, academic, partner-generated, and real-world datasets. This richer context enables AI to tackle harder scientific questions — target selection, toxicity prediction, biomarker discovery, and patient stratification — that internal datasets alone rarely support.

Privacy-preserving frameworks like federated learning are emerging to enable multi-organization collaboration without sharing raw data, a key unlock for regulated environments.

This marks a shift from closed, siloed R&D to networked, data-rich ecosystems. Organizations are now reckoning with how to integrate this data securely, consistently, and in a way compliant with IP and regulatory requirements.

> "Knowledge management wasn't so much a key quality system enabler as it was a heavy lift. Now knowledge just keeps pace with our questions. It's startling to see how quickly this became normal.
>
> Chris Bell
> VP Quality Assurance
> prime medicine

**Internal and external data pipelines**

**YoY growth of external data**

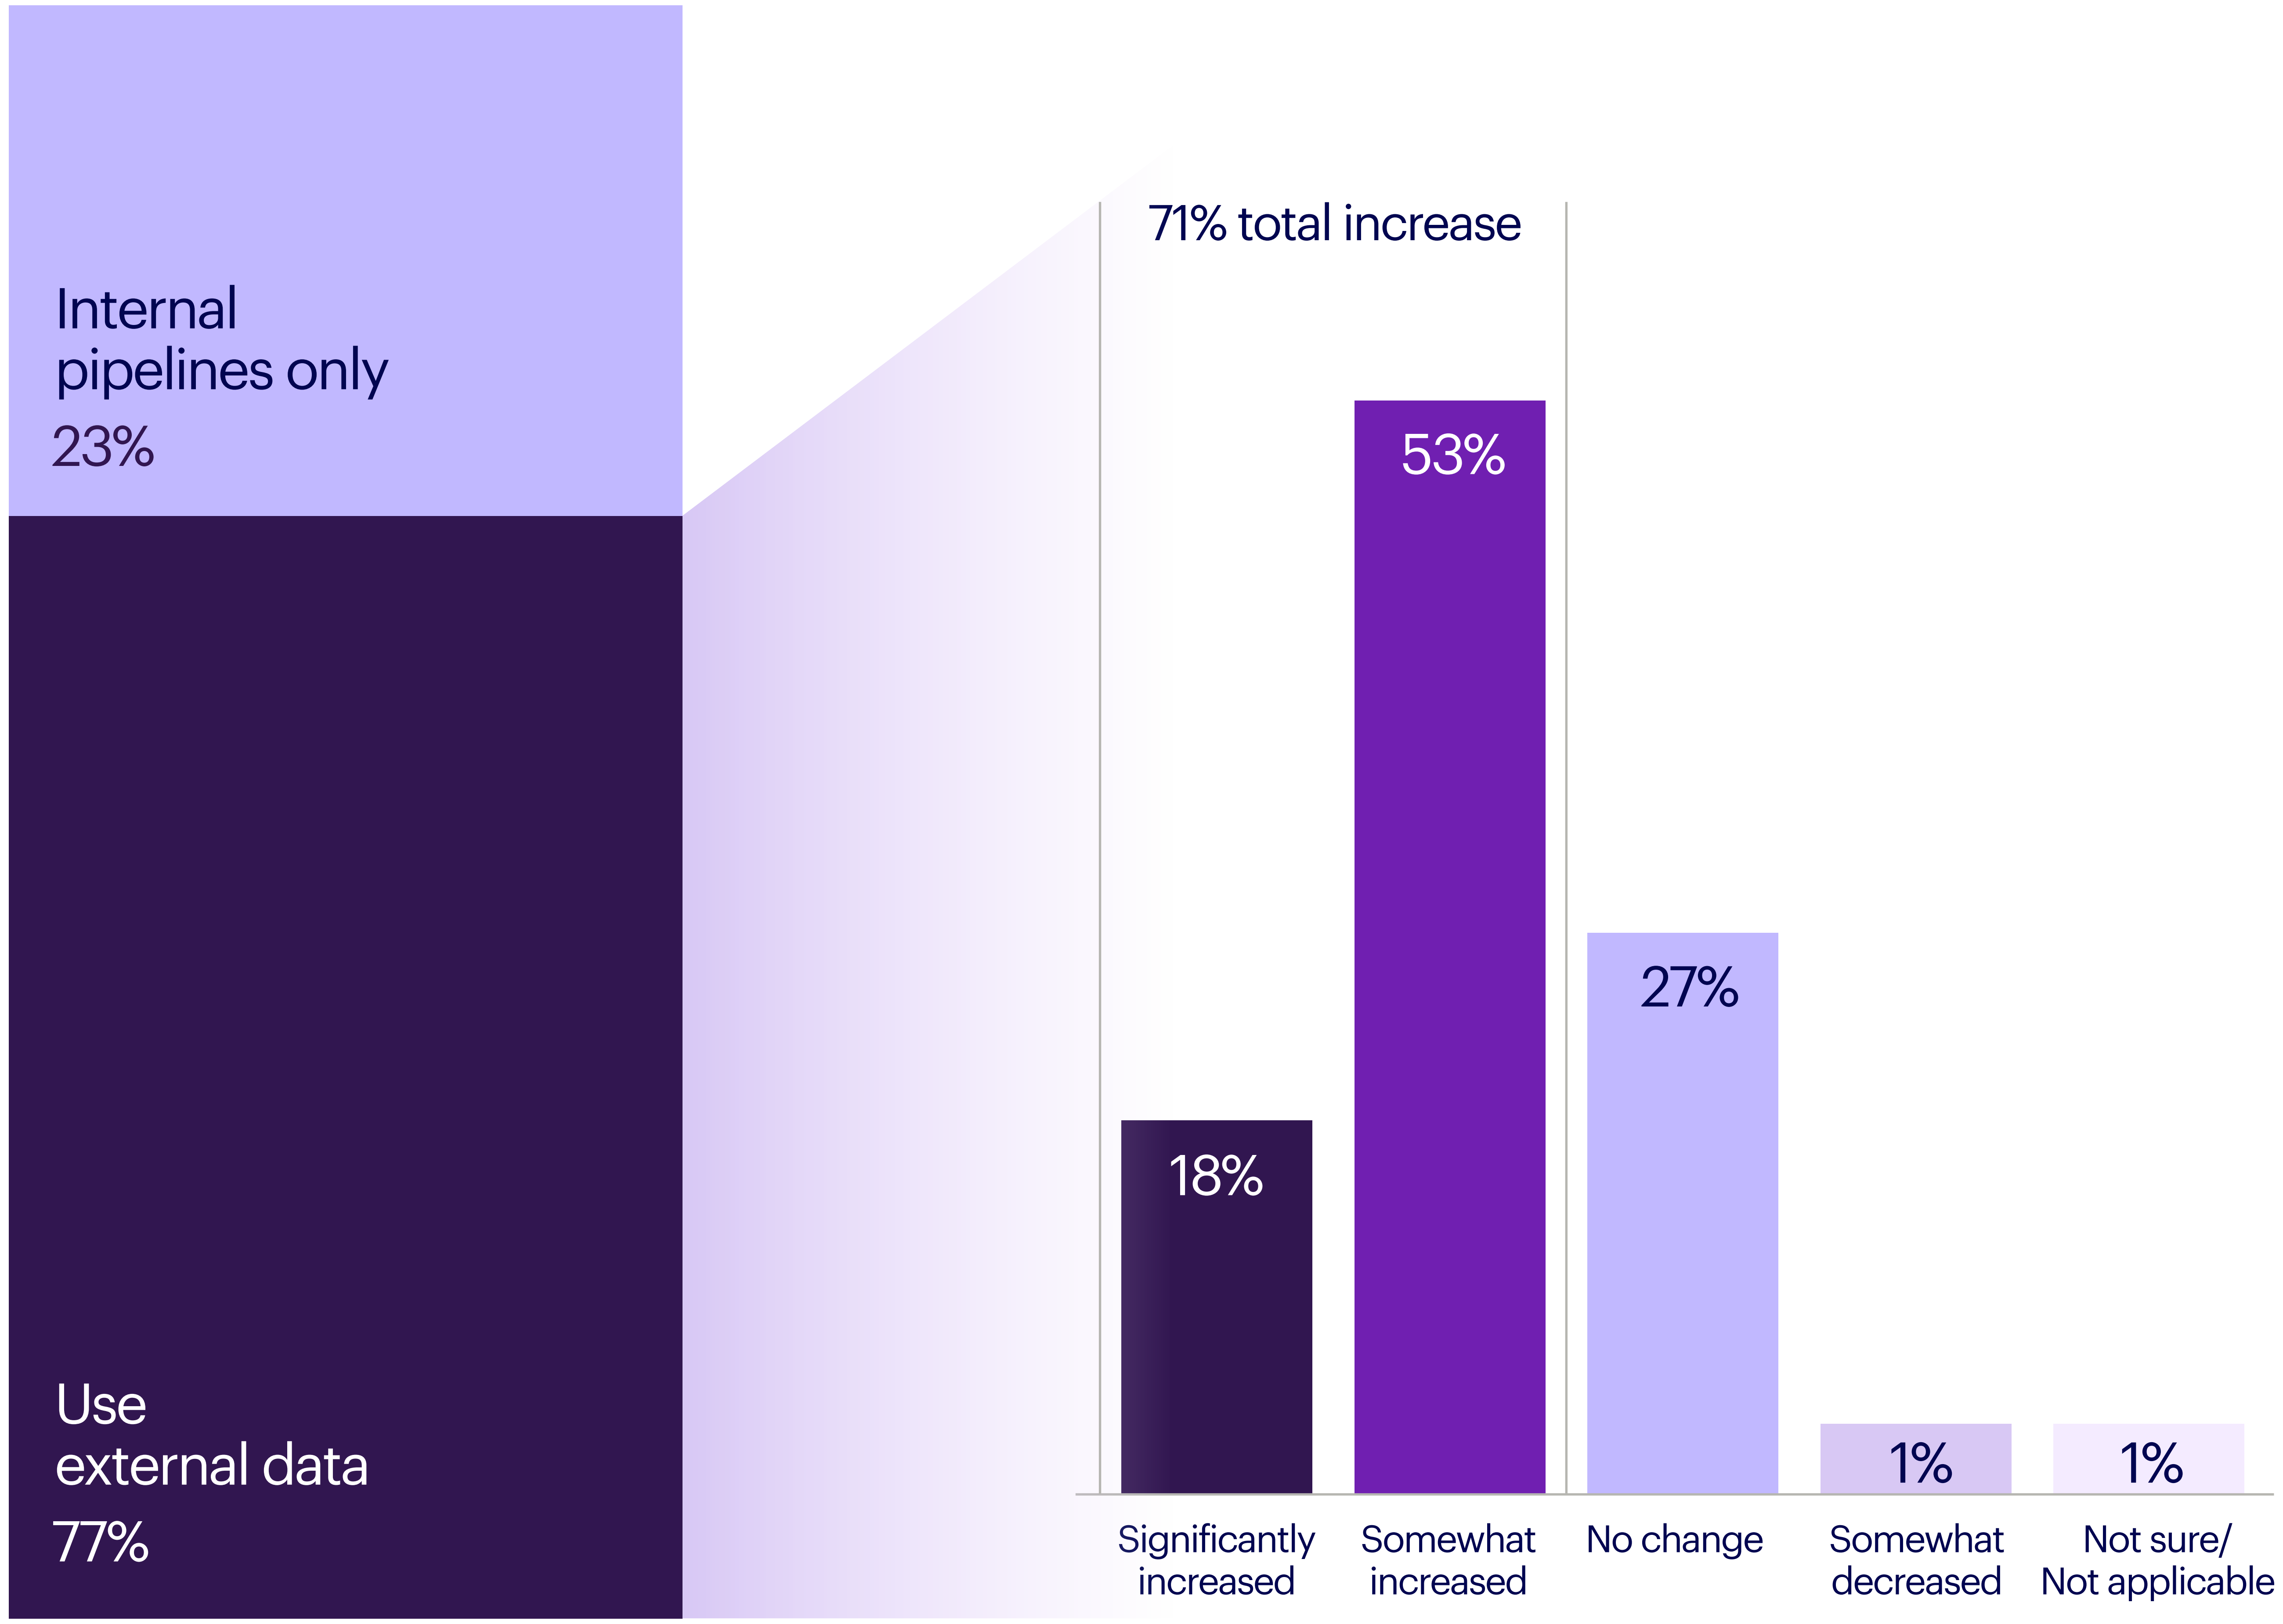


N = 104

N = 73

# Biotech tap diverse external data sources, with clear differences by company size

Small biotech overwhelmingly lean on freely available public datasets (90%), while large biotech are expanding into licensed commercial databases, real-world evidence, industry consortium data, and synthetic datasets.

These different sourcing strategies signal a growing divide: smaller teams optimize for access and speed, while larger organizations invest in breadth, depth, and exclusivity of data.

> "For many biotechs, the promise of AI-enabled drug discovery remains unrealized due to a fundamental hurdle: they simply don't have access to the large-scale, high-quality data needed to impact decisions and train truly effective models. Lilly TuneLab was created to be an equalizer, using federated learning to provide access to select models trained on decades of Lilly research data. This helps propel biotech innovation and lets models learn across organizational boundaries, while data stays sovereign.

Aliza Apple
Global Head of Lilly TuneLab
and Catalyze 360 AI/ML Eli Lilly

**Sources of external data for AI-driven R&D**

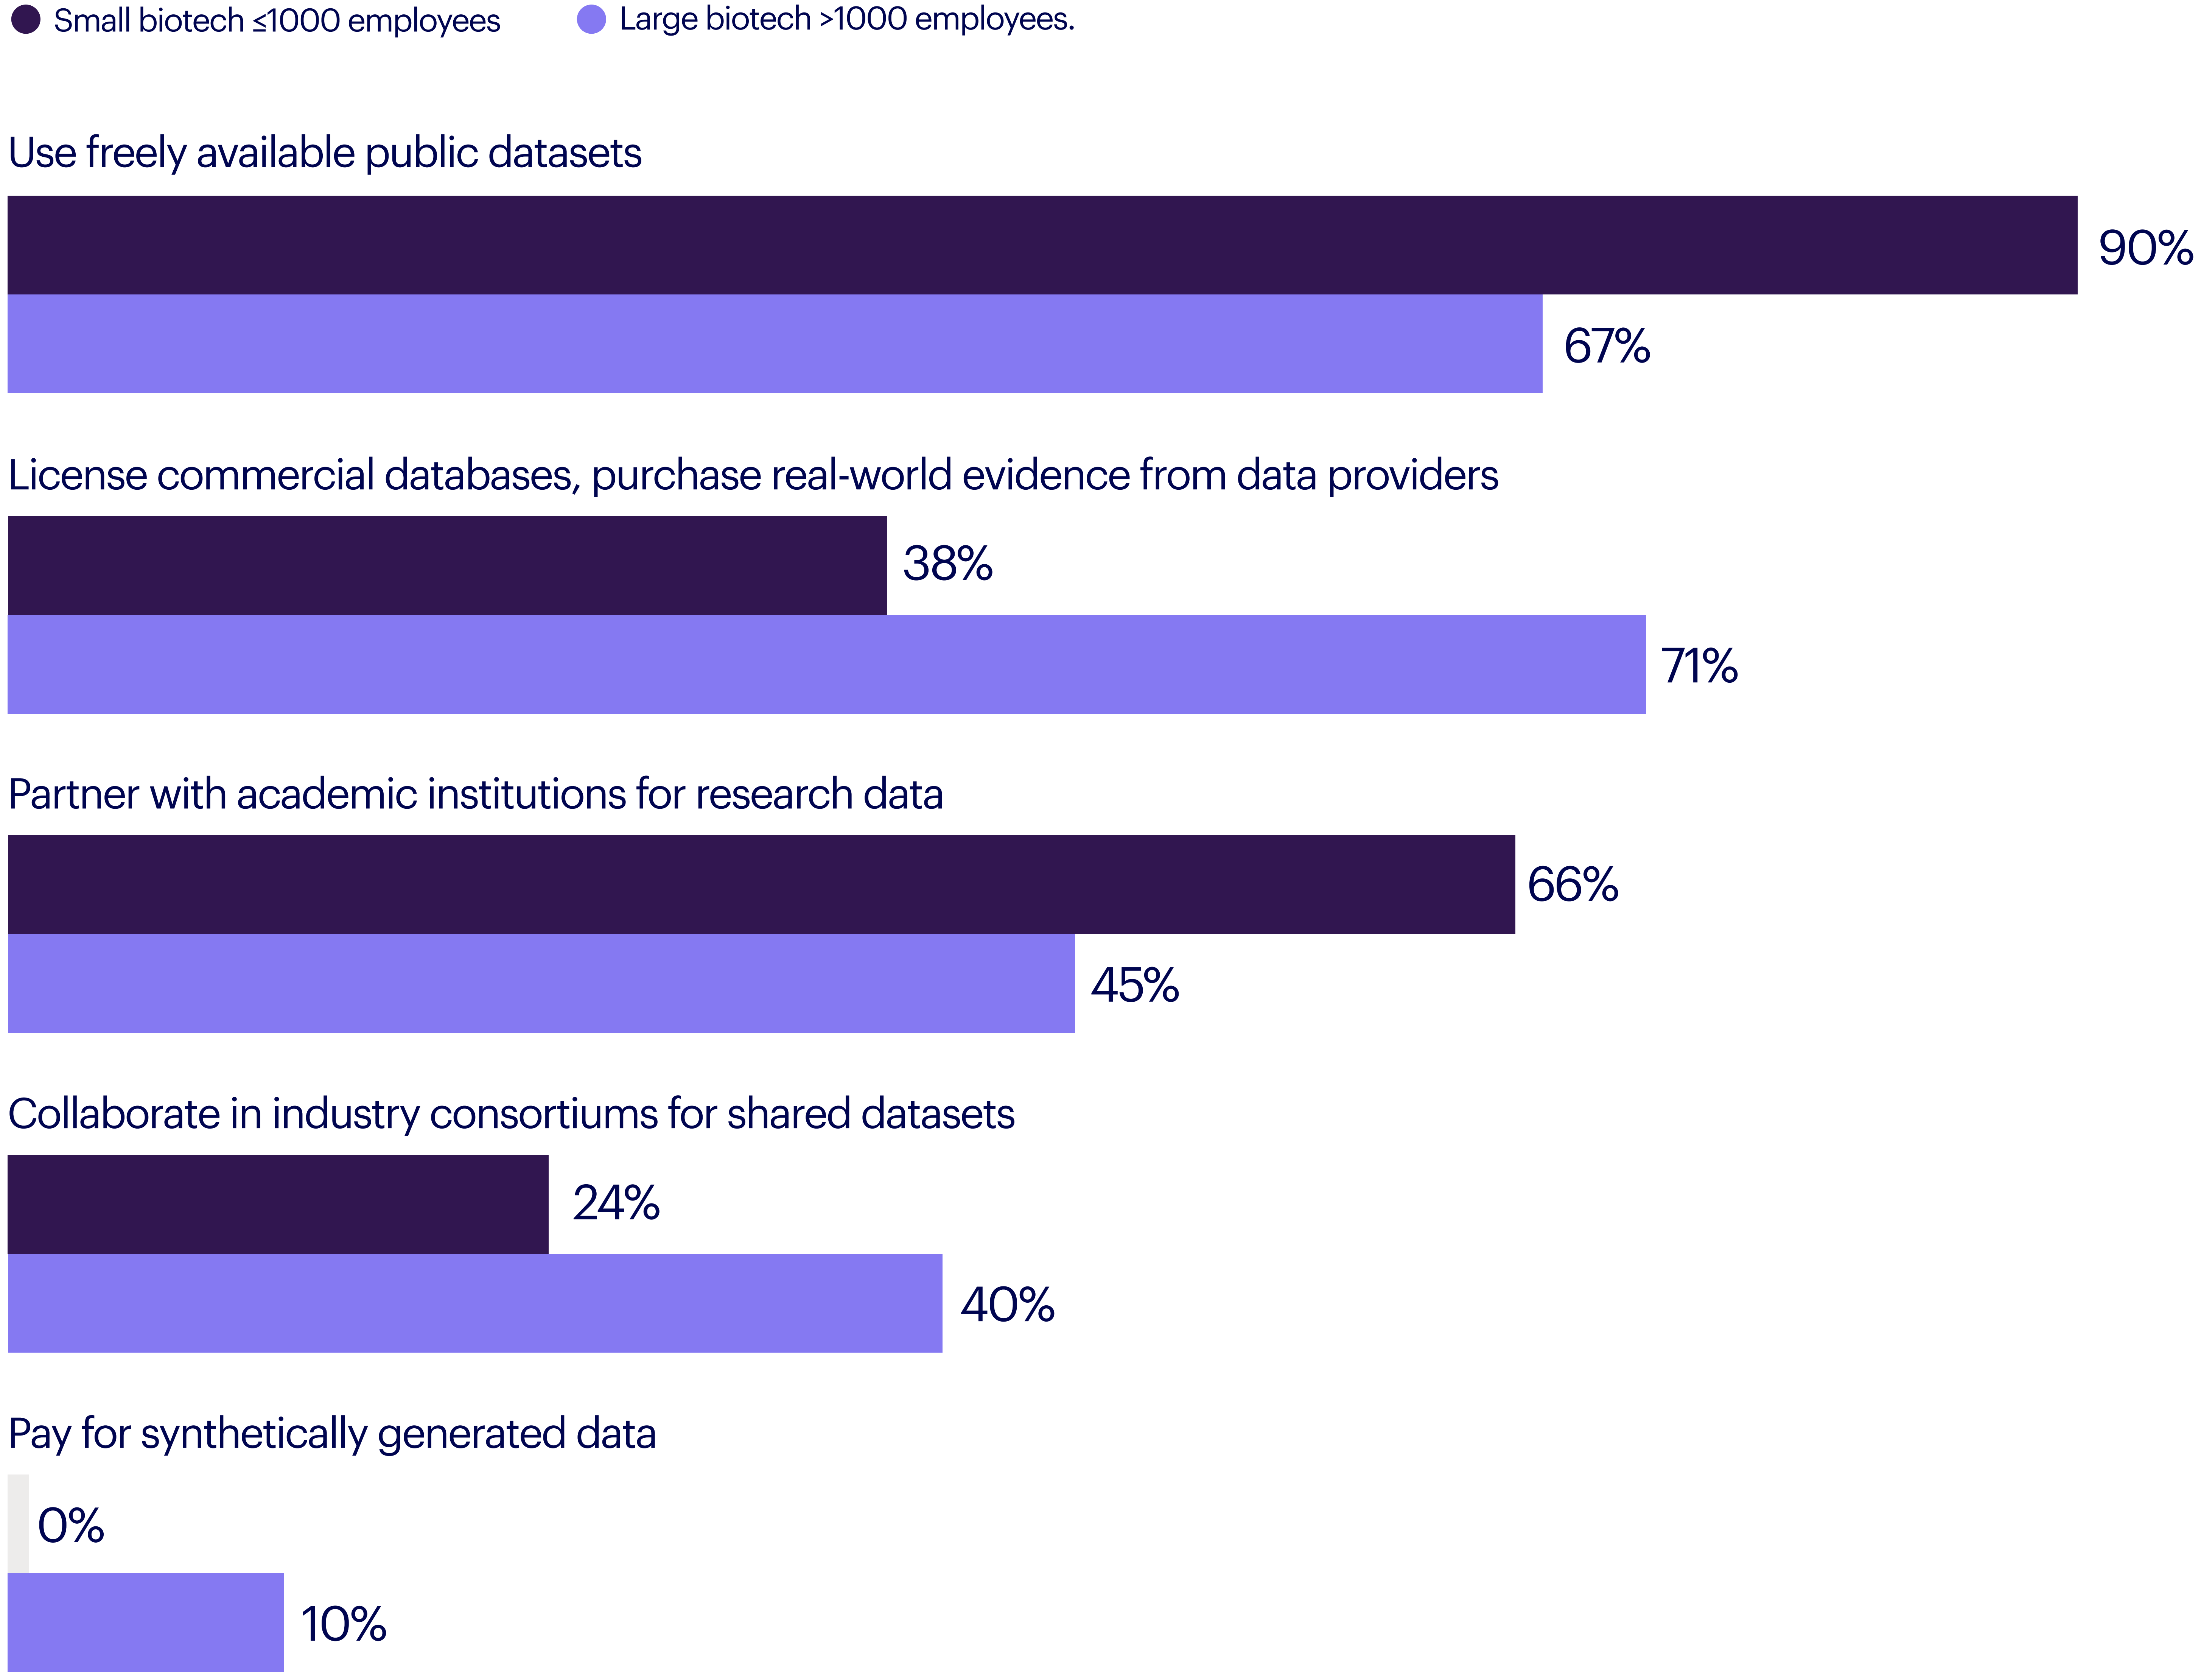


N = 32 for small biotech, N = 42 for large biotech
% indicates adoption

# AI is now a first step in data inquiry

AI tools are rapidly becoming the starting point for scientific inquiry, shifting research from manual search to instant insight.

**How scientists query external data for R&D**

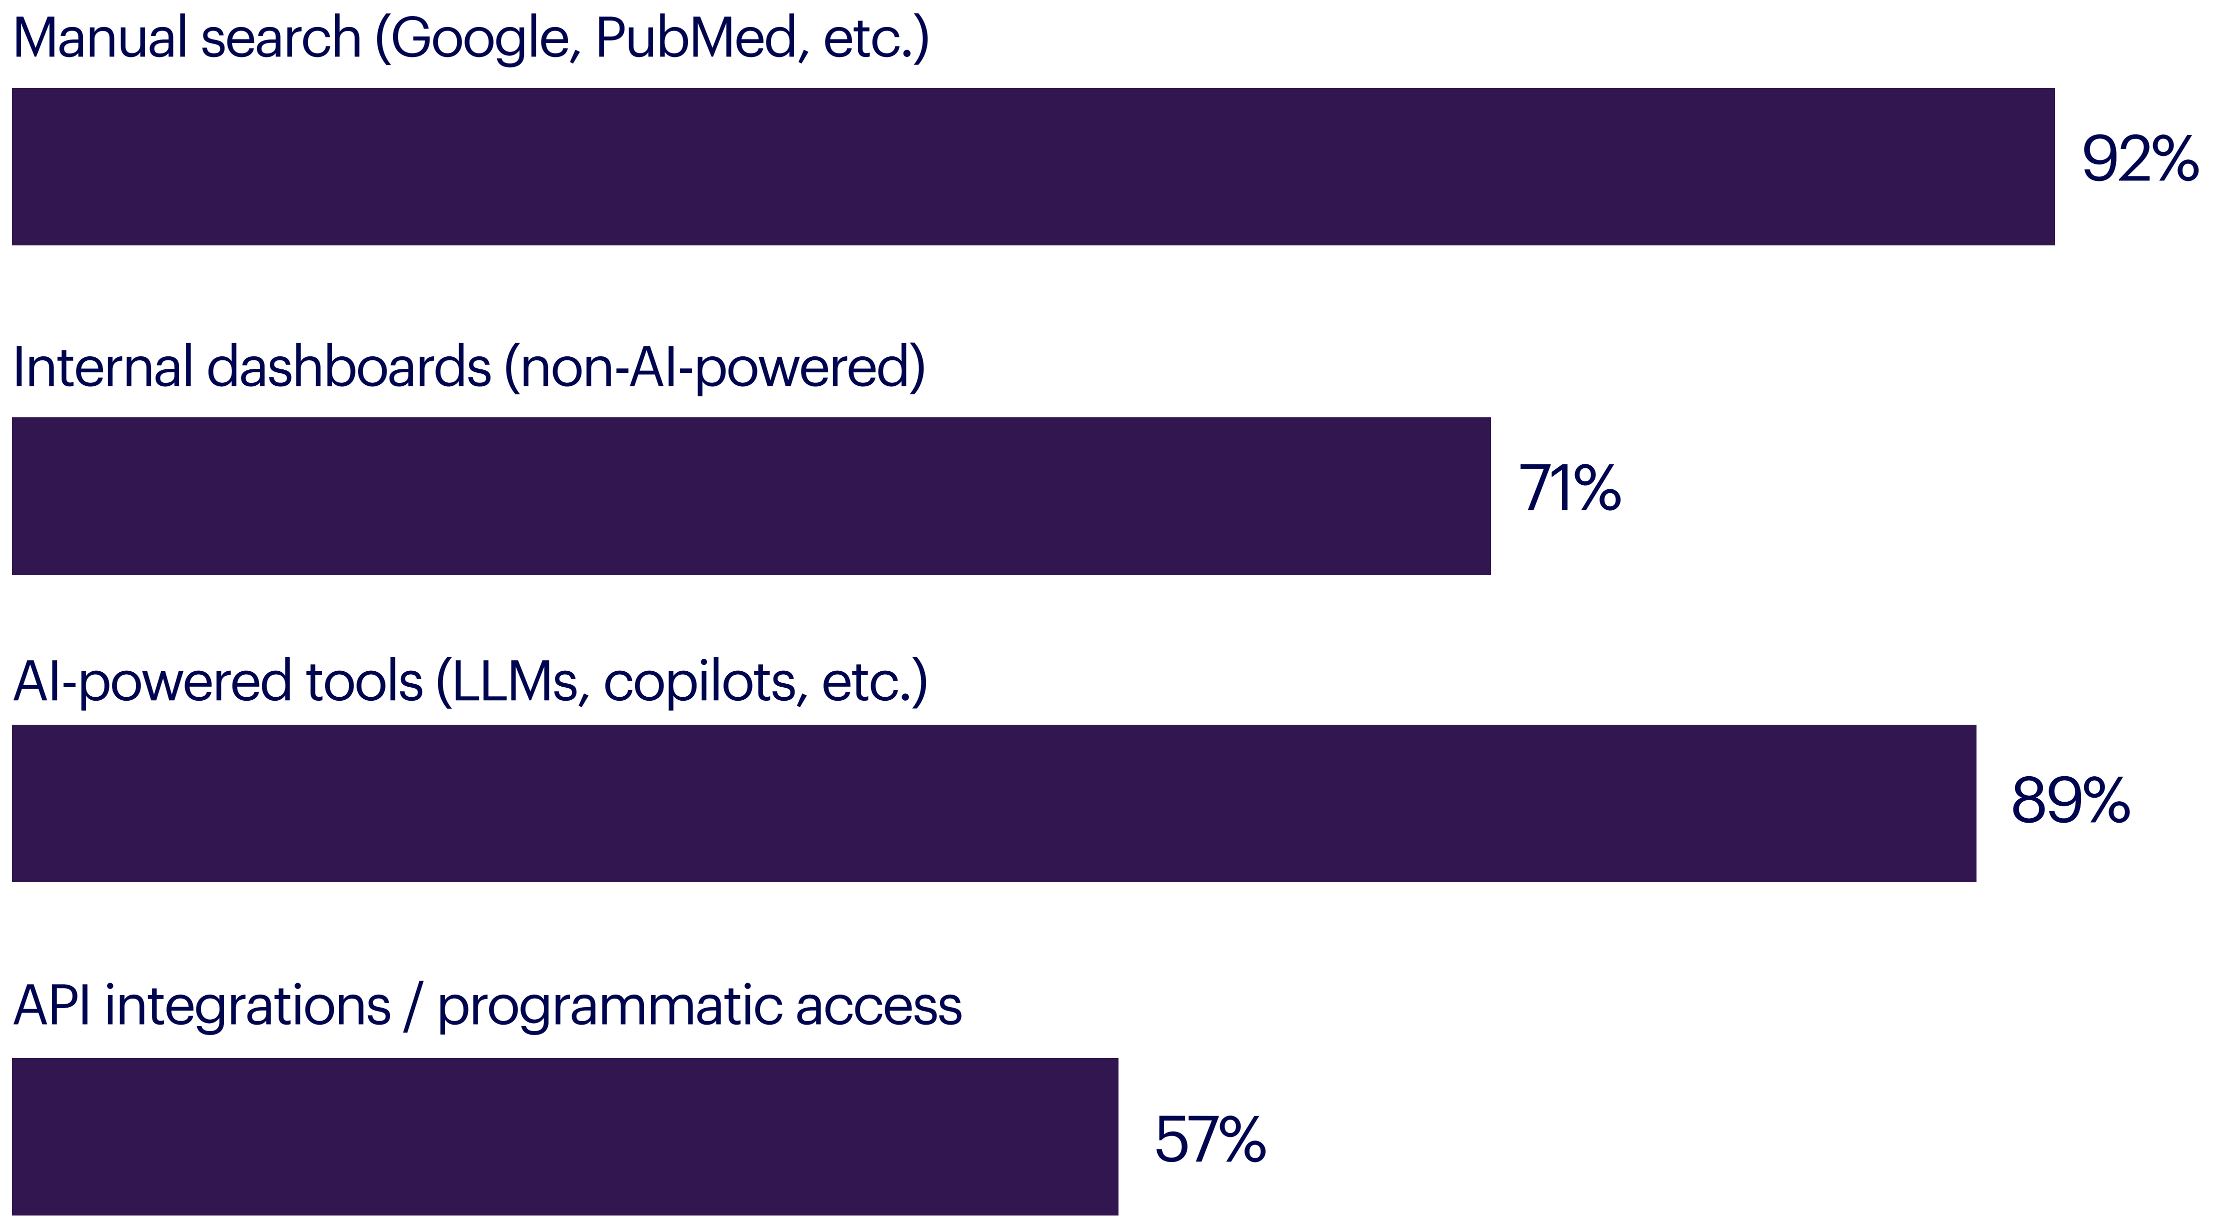


N = 104, % represents adoption

**Percent reporting increase in querying external data (YoY)**

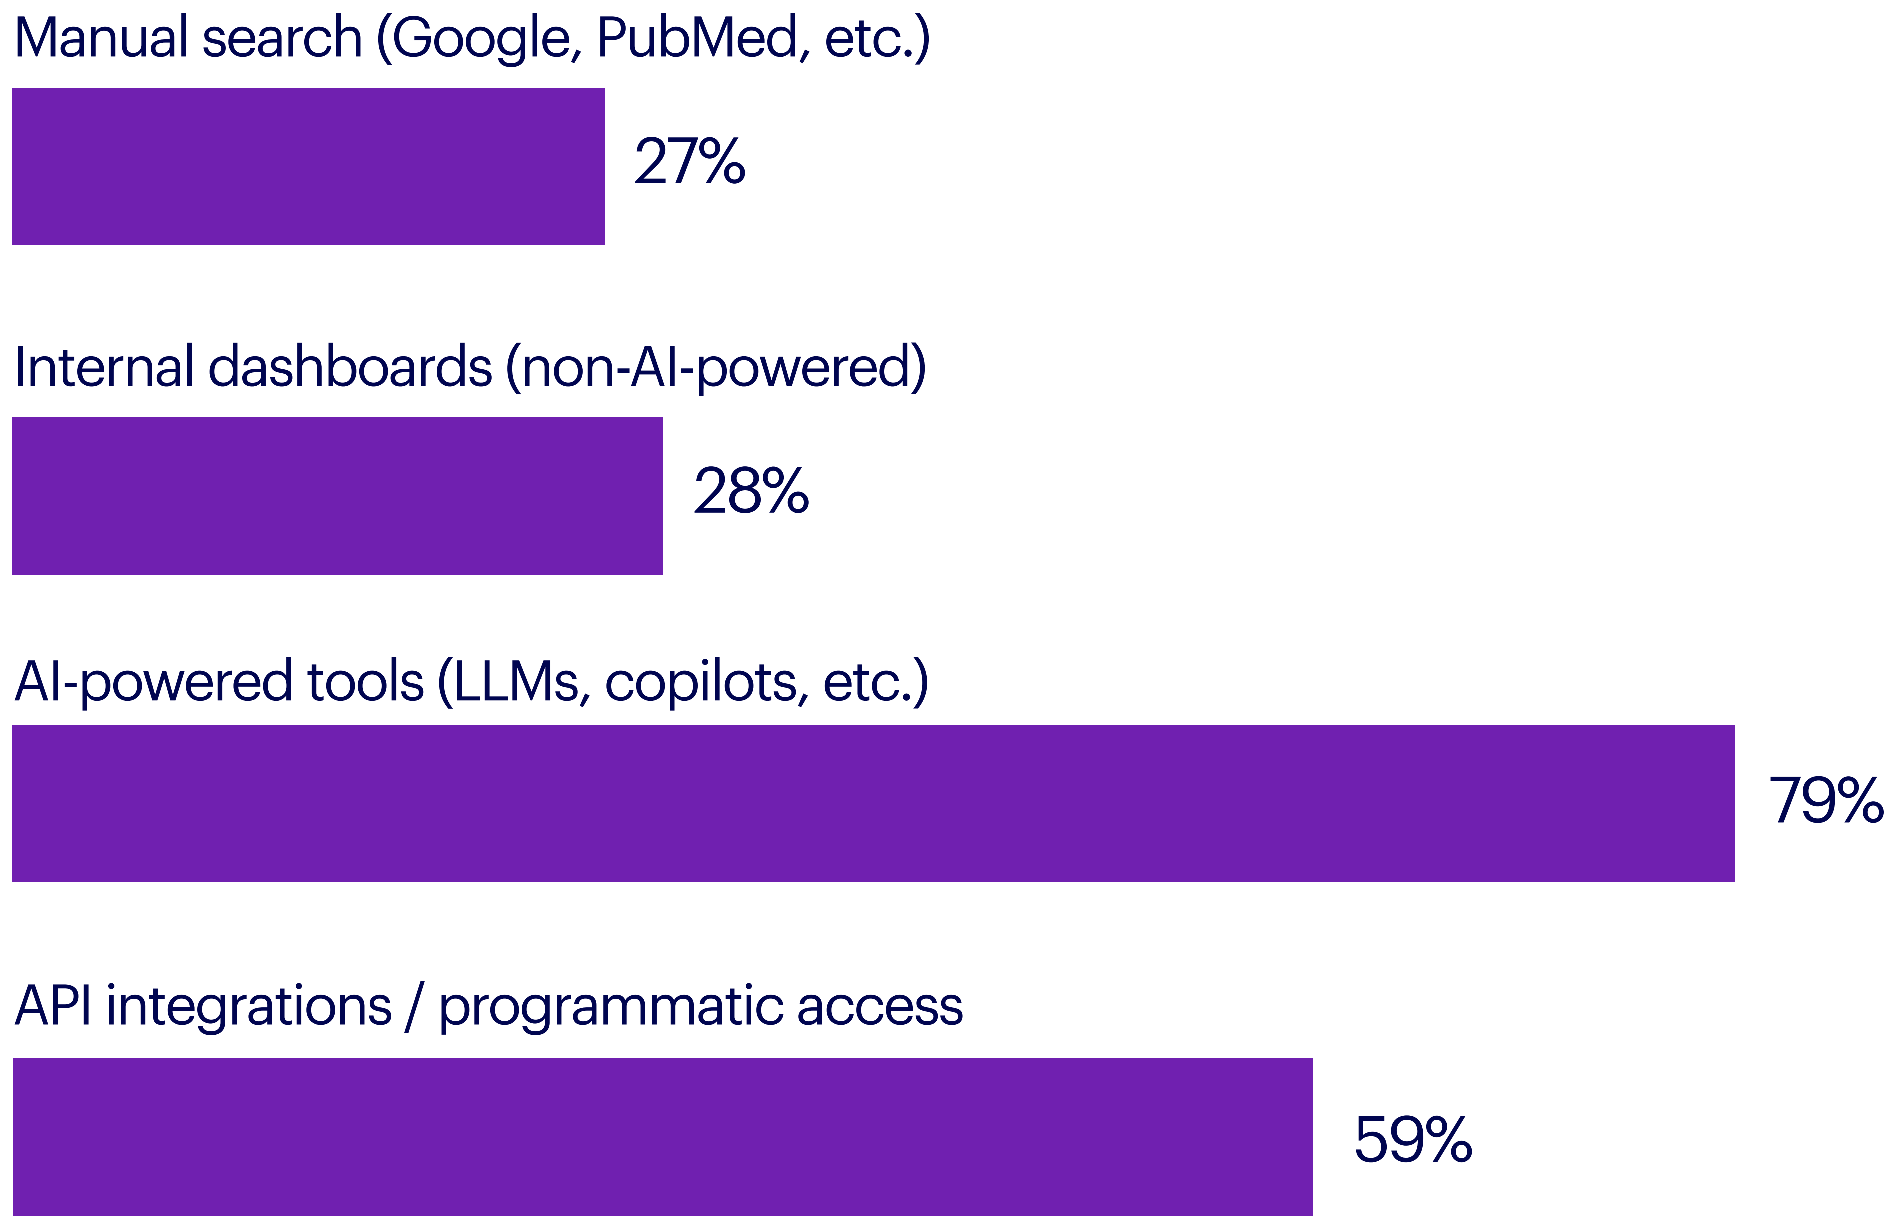


89% of scientists now use AI tools like LLMs and copilots to interrogate external data, nearly matching manual search (92% citing) and growing much faster year over year. Instead of reading through hundreds of papers themselves, scientists increasingly begin with AI tools that can find, filter, and synthesize the evidence for them — summarizing results, highlighting patterns, and connecting insights across the literature. What began with general-purpose models is evolving into specialized “deep research” agents built for scientific workflows, making AI the emerging first step in data inquiry.

# AI maturity mirrors data maturity

Most biotech describe their data foundations as “developing” with about half sitting in the middle of the maturity curve. Their data works well enough for individual teams, but it’s too fragmented for AI to operate across workflows or experiment types. Critical details — experimental context, batch metadata, process history — often never make it into the system. This leads to models that only perform on the dataset they were trained on, or teams rerunning experiments because earlier results can’t be trusted or reused.

Organizations that use AI more regularly look different. They’ve invested in stronger data foundations: 39% of high AI adopters rate their infrastructure as highly integrated, compared to just 25% of low adopters.

In contrast, 31% of low AI adopters say their systems are still siloed or early-stage, while only 9% of high adopters say the same.

> “The models are getting much better at R&D-related tasks, to the point that for many workflows, the bottleneck is in the product layer. Most organizations are still running workflows across dozens of product surfaces, with data fragmented across different locations. The next evolution is a unified interface for AI, connecting continuous model improvements to the experimental data that drives R&D.

Eric Kauderer-Abrams
Head of Life Sciences

ANTHROPIC

## Current levels of data integration for AI in R&D

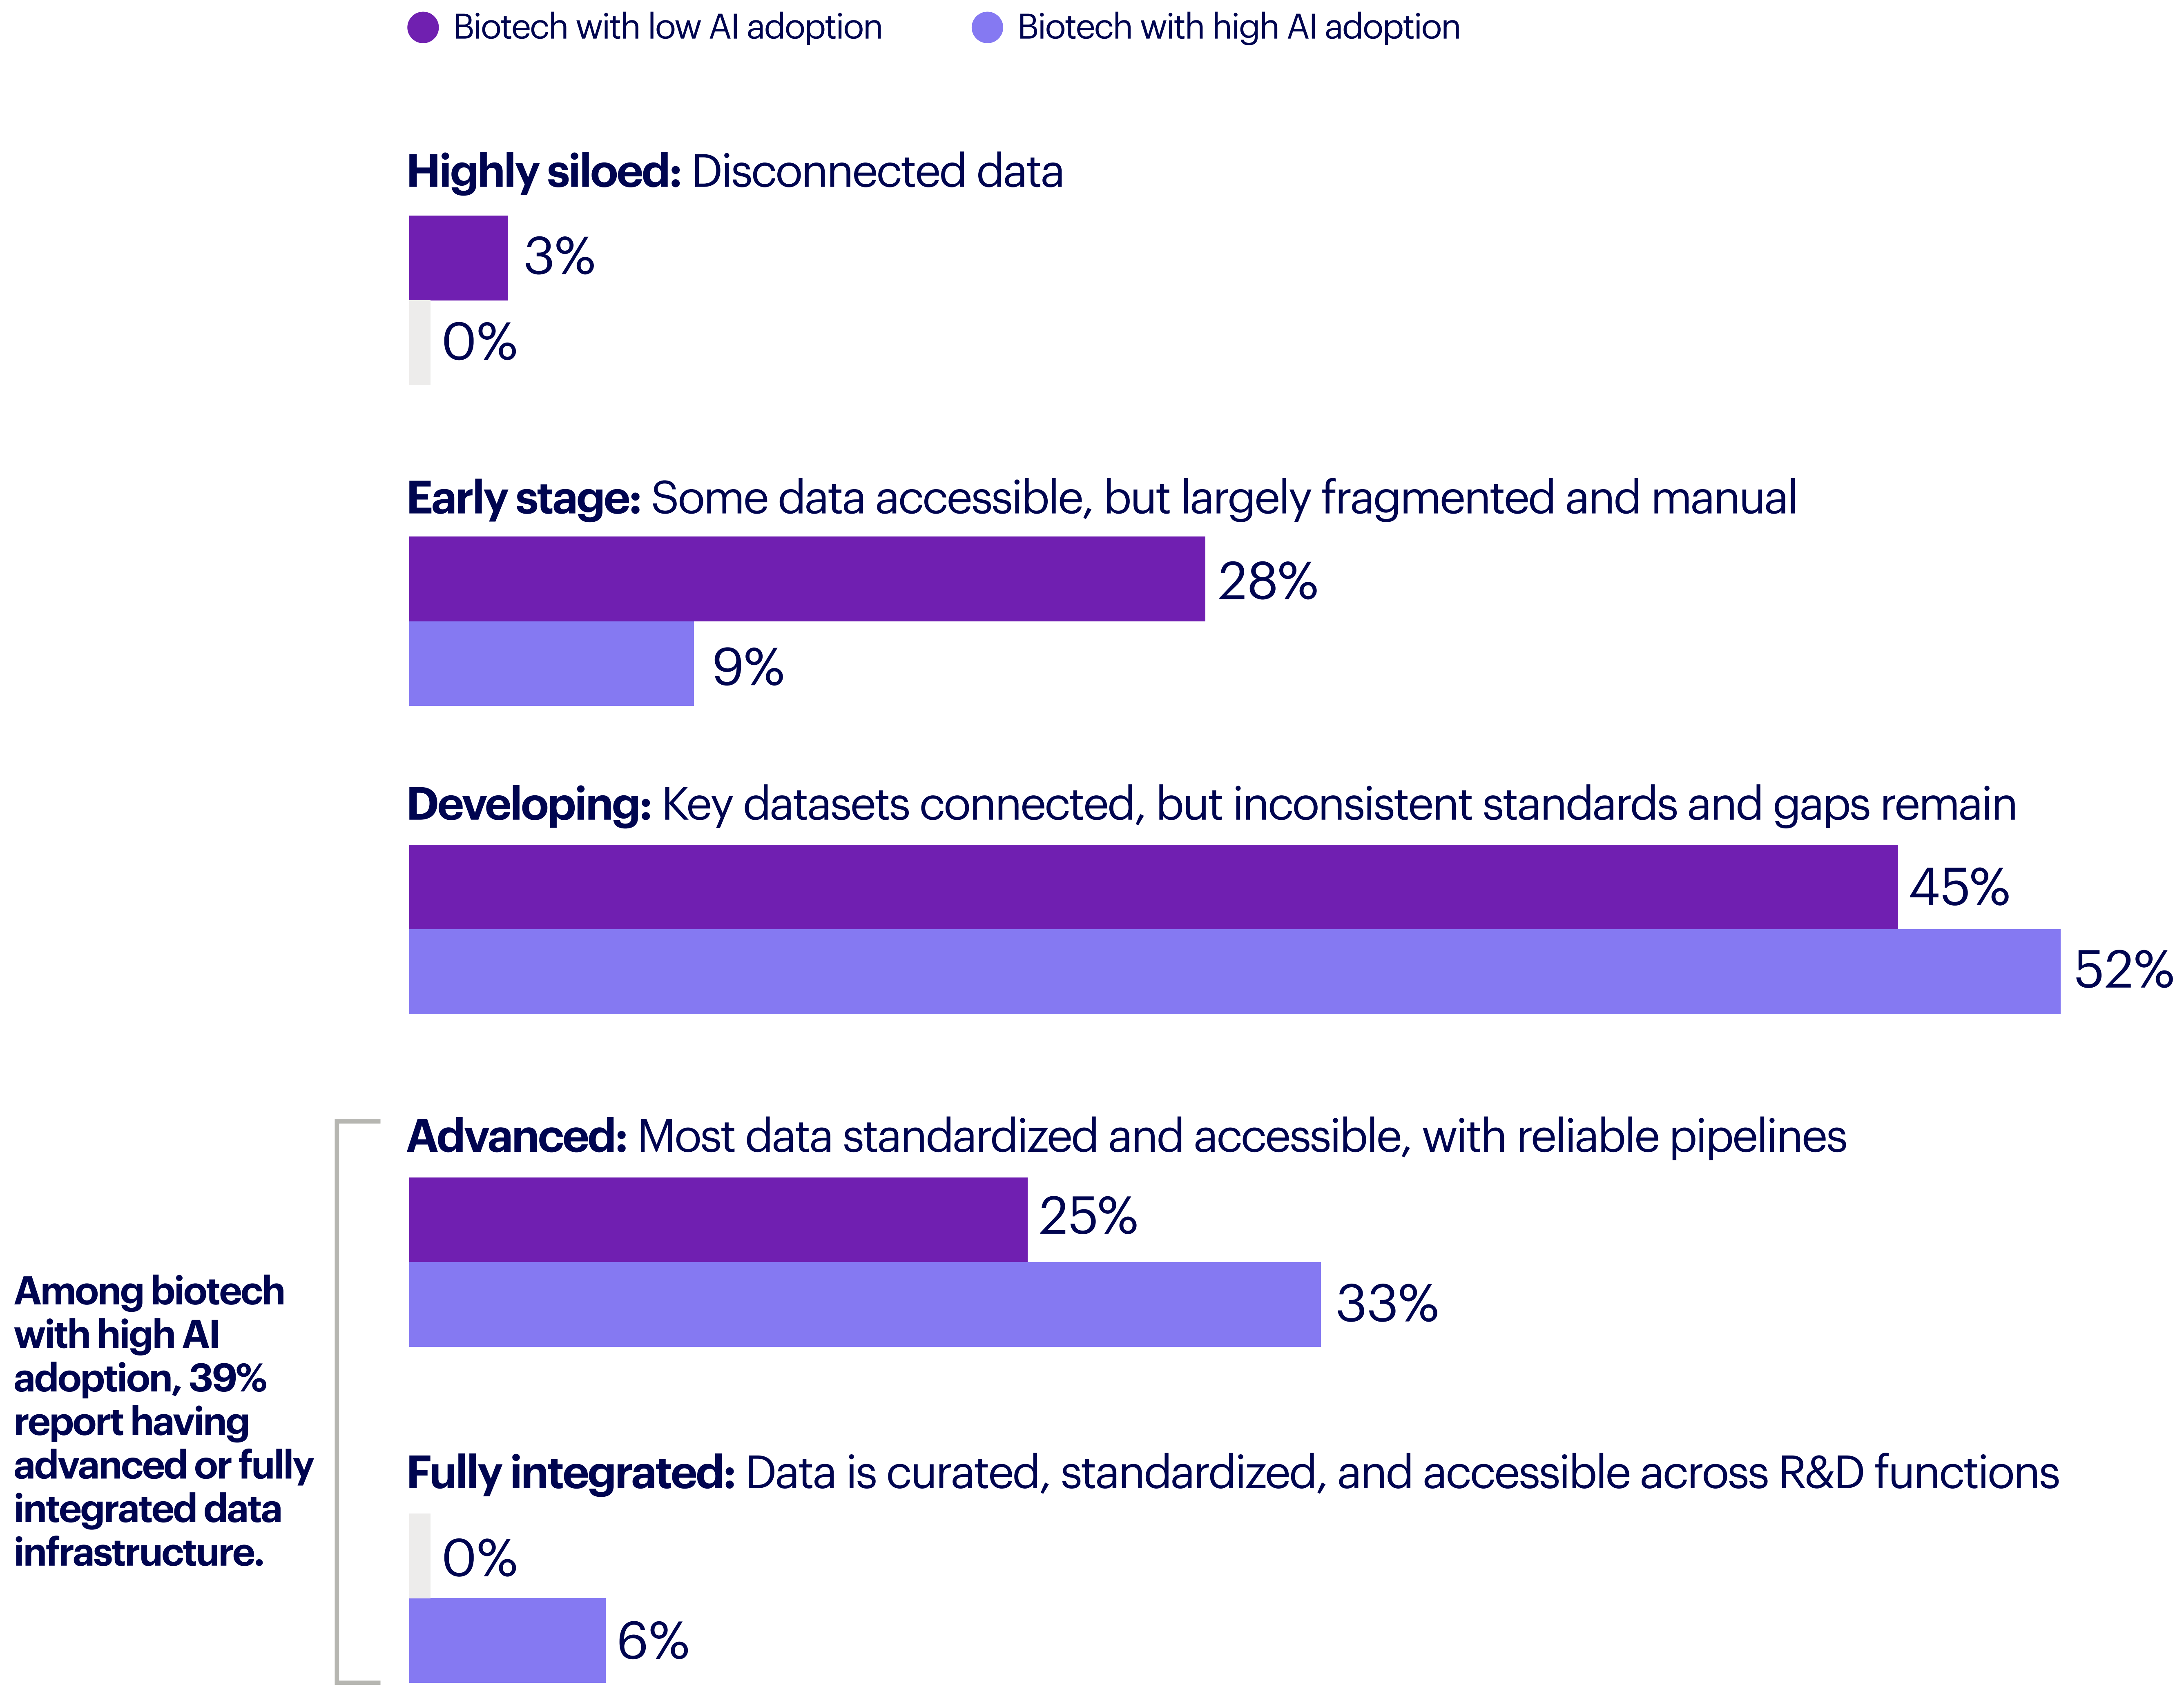


Biotech with high AI adoption N = 64, Biotech with low AI adoption N = 40, as determined by responding use AI regularly (high AI adoption) versus use AI in highly specific use cases (low AI adoption)

Numbers do not add to 100 due to rounding

# Leaders focus on wet-dry lab integration

## Wet-dry lab integration to support AI in R&D

- Not integrated: Systems operate separately with little to no data flow
- Minimally integrated: Limited connections, mostly manual data transfer
- Partially integrated: Some automated data flow, but significant gaps remain
- Well integrated: Most systems connected with reliable data exchange
- Fully integrated: Experimental and computational data flow seamlessly between systems

### Biotech with low AI adoption

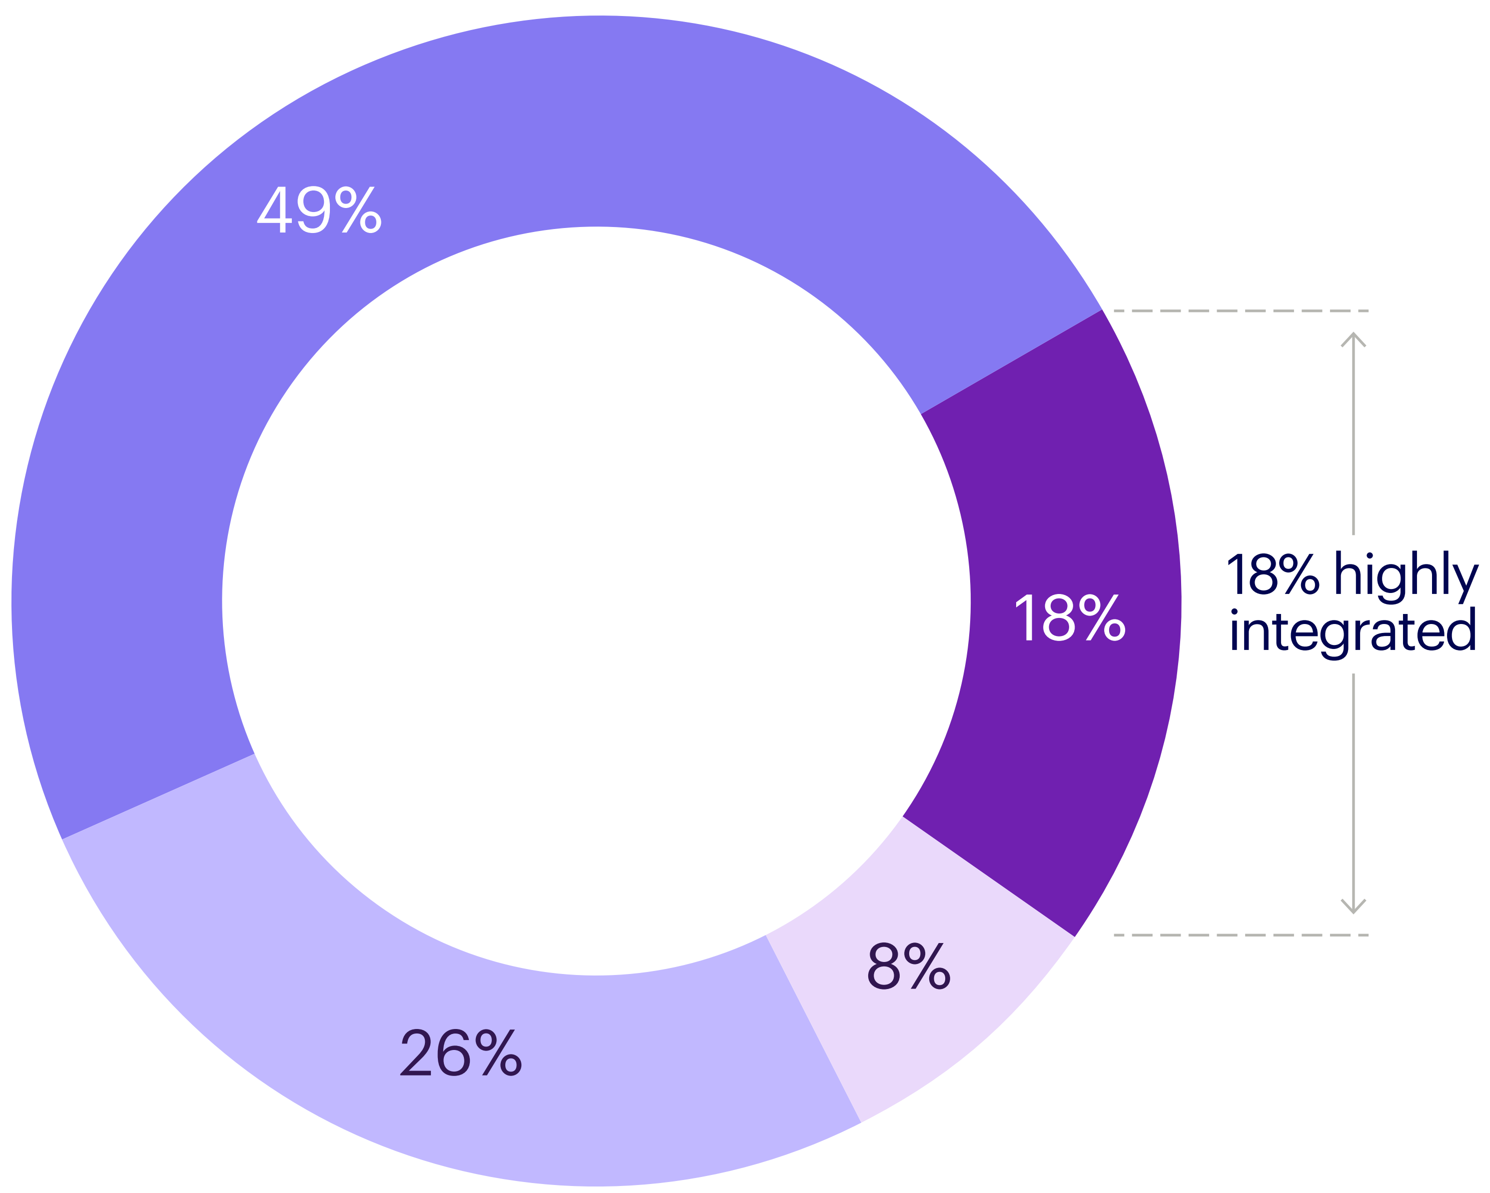


### Biotech with high AI adoption

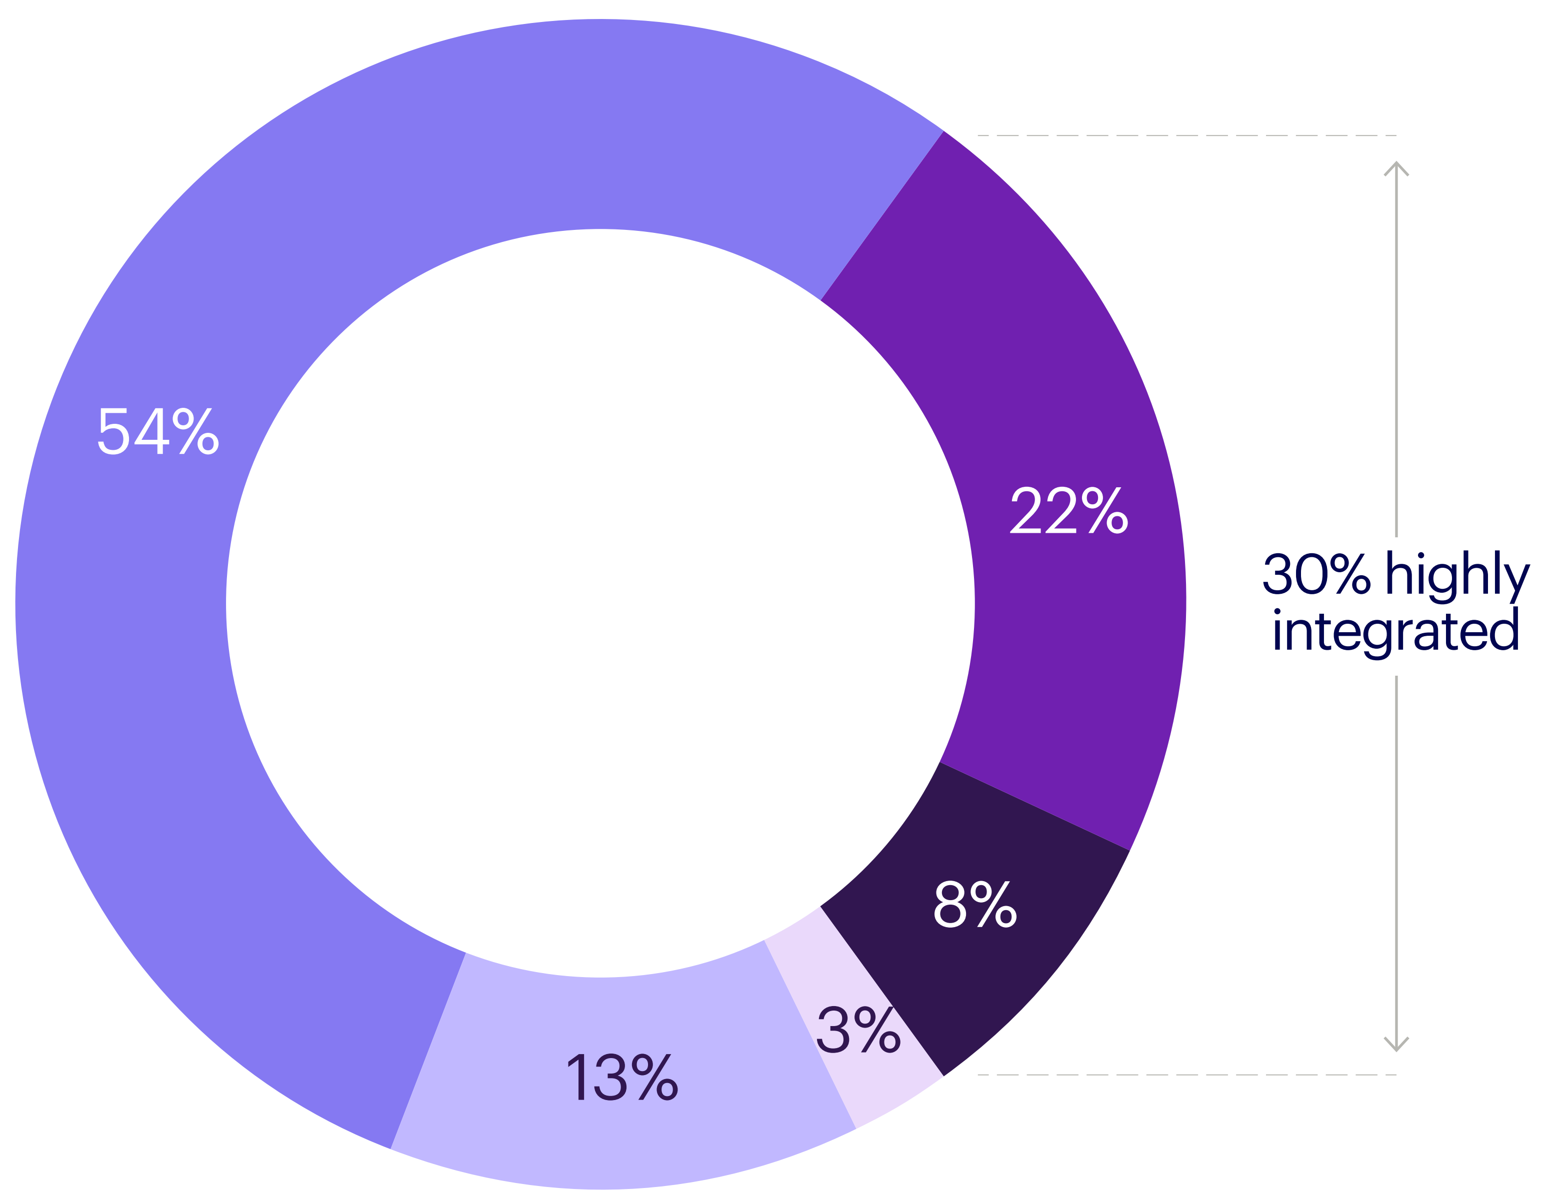


Biotech with high AI adoption N = 64, Biotech with low AI adoption N = 40, as determined by responding use AI regularly (high AI adoption) versus use AI in highly specific use cases (low AI adoption)

Numbers do not add to 100 due to rounding

Organizations using AI regularly are nearly 2x more likely to report strong wet–dry lab integration (30% vs. 18% citing).

Improving wet–dry lab integration means capturing data automatically at the bench and closing the loop between computation and experiment. Today, that loop is often broken: a wet-lab scientist might run an assay and wait days or weeks for a model to analyze the results. By the time the insight arrives, it's too late to adjust the next experiment. And when results show up manually, out of context, or with missing metadata, AI can't help; it just adds noise to an already noisy process.

Leaders are solving this by building closed-loop R&D systems that link simulation, experiment design, automated data capture, and rapid computational feedback. The goal is that scientists act on AI-generated insights in real time.

# Quality and interoperability are top priorities for improving data for AI

As labs move toward more automation and closed-loop experimentation, the bar for data gets much higher. AI can only operate end-to-end if the underlying data is consistent, trusted, and able to move cleanly across teams and systems. That means capturing data in structured, machine-readable formats from the moment it's created.

When asked what matters most with improving data for AI, leaders point first to data quality, standardization, and accuracy. Next comes data interoperability, followed by tighter wet–dry lab integration, and data governance.

> "The race isn't to the best model; it's to the best data-to-decision loop. As co-scientist agents mature and multimodal foundation models unify sequence, structure, assay, and language into a single reasoning layer, we're witnessing the emergence of AI as the operating system for drug discovery, with chat as the universal interface. The winners won't be those with the most sophisticated algorithms, but those who've built the end-to-end infrastructure where data flows from bench to model to decision and back, continuously, autonomously, and at the speed of science. This is how the next generation of medicines gets made.

Yves Fomekong Nanfack
Head of AI/ML Research
Takeda

### Top data priorities for AI-driven R&D

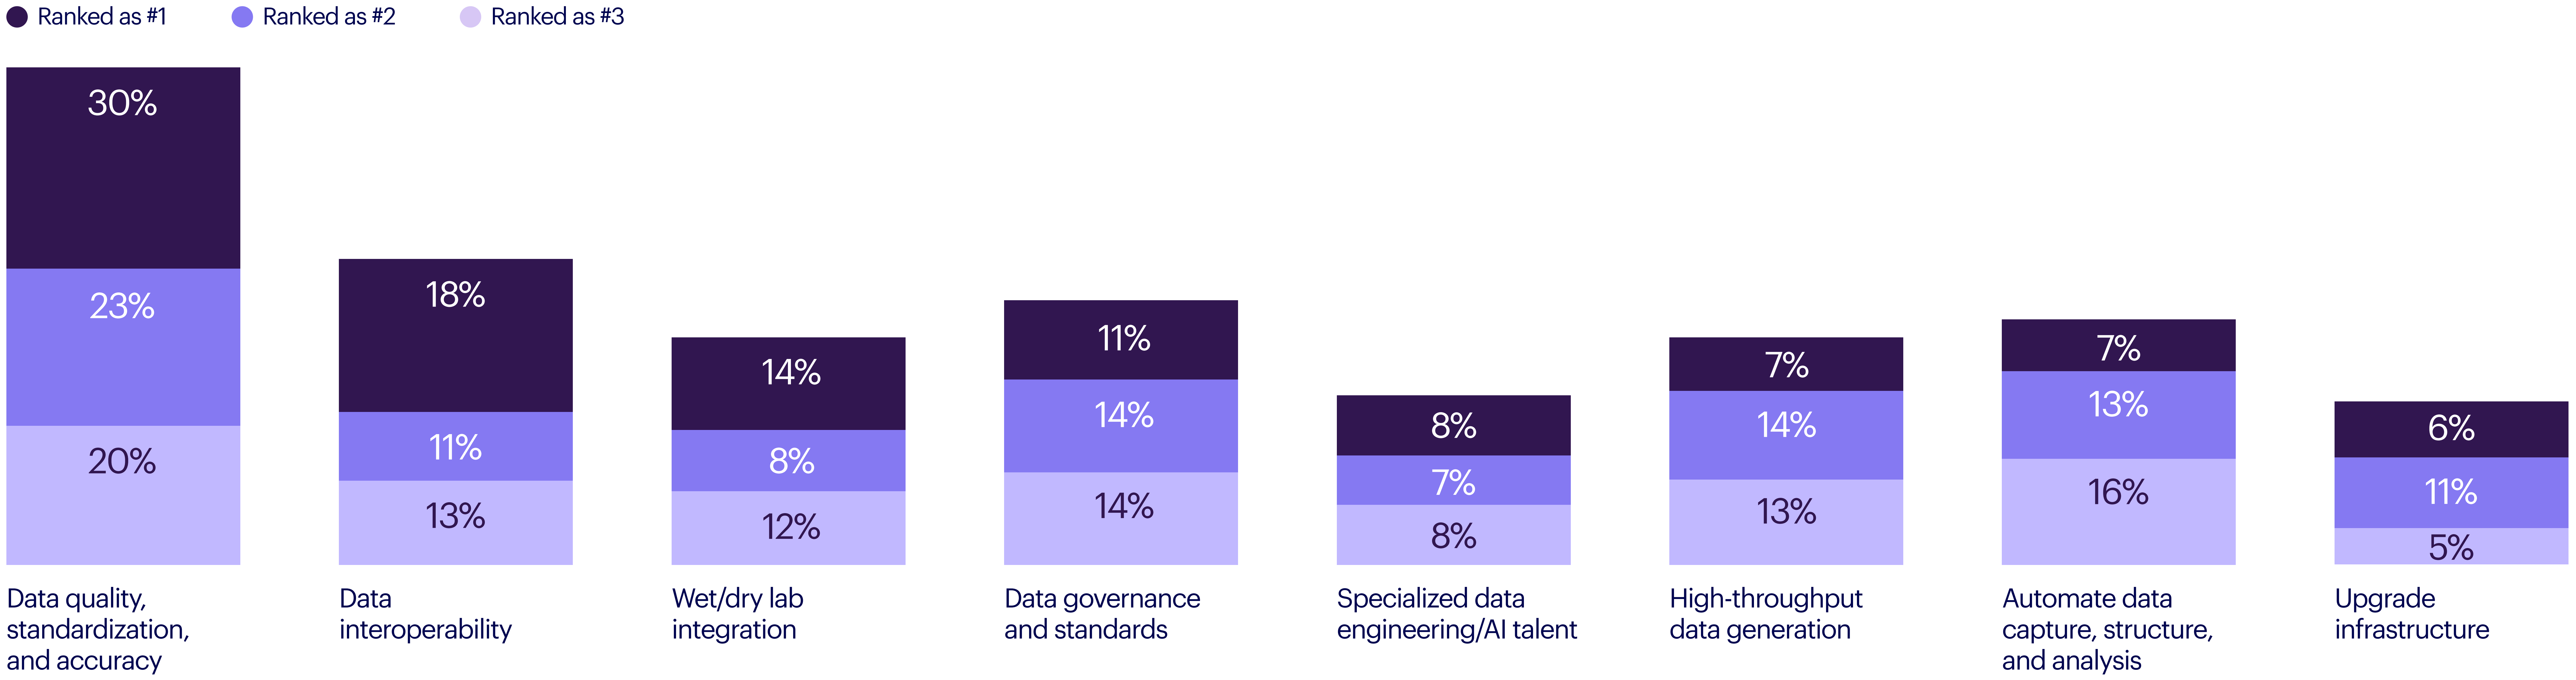


N = 104

# Talent & Org Structure

## AI is reshaping the org chart, starting with the bench

Biotechs are organizing their work so scientists and technologists can collaborate more closely in small, fast-moving groups. Companies are developing this talent internally, creating hybrid roles — people who understand both the science and the AI needed to support it. This shift is helping teams connect models, data, and experiments more directly, and is becoming a common pattern in how biotechs structure their AI work.

# New R&D org chart: Hybrid teams built around speed, context, and shared language

### How AI leadership is structured

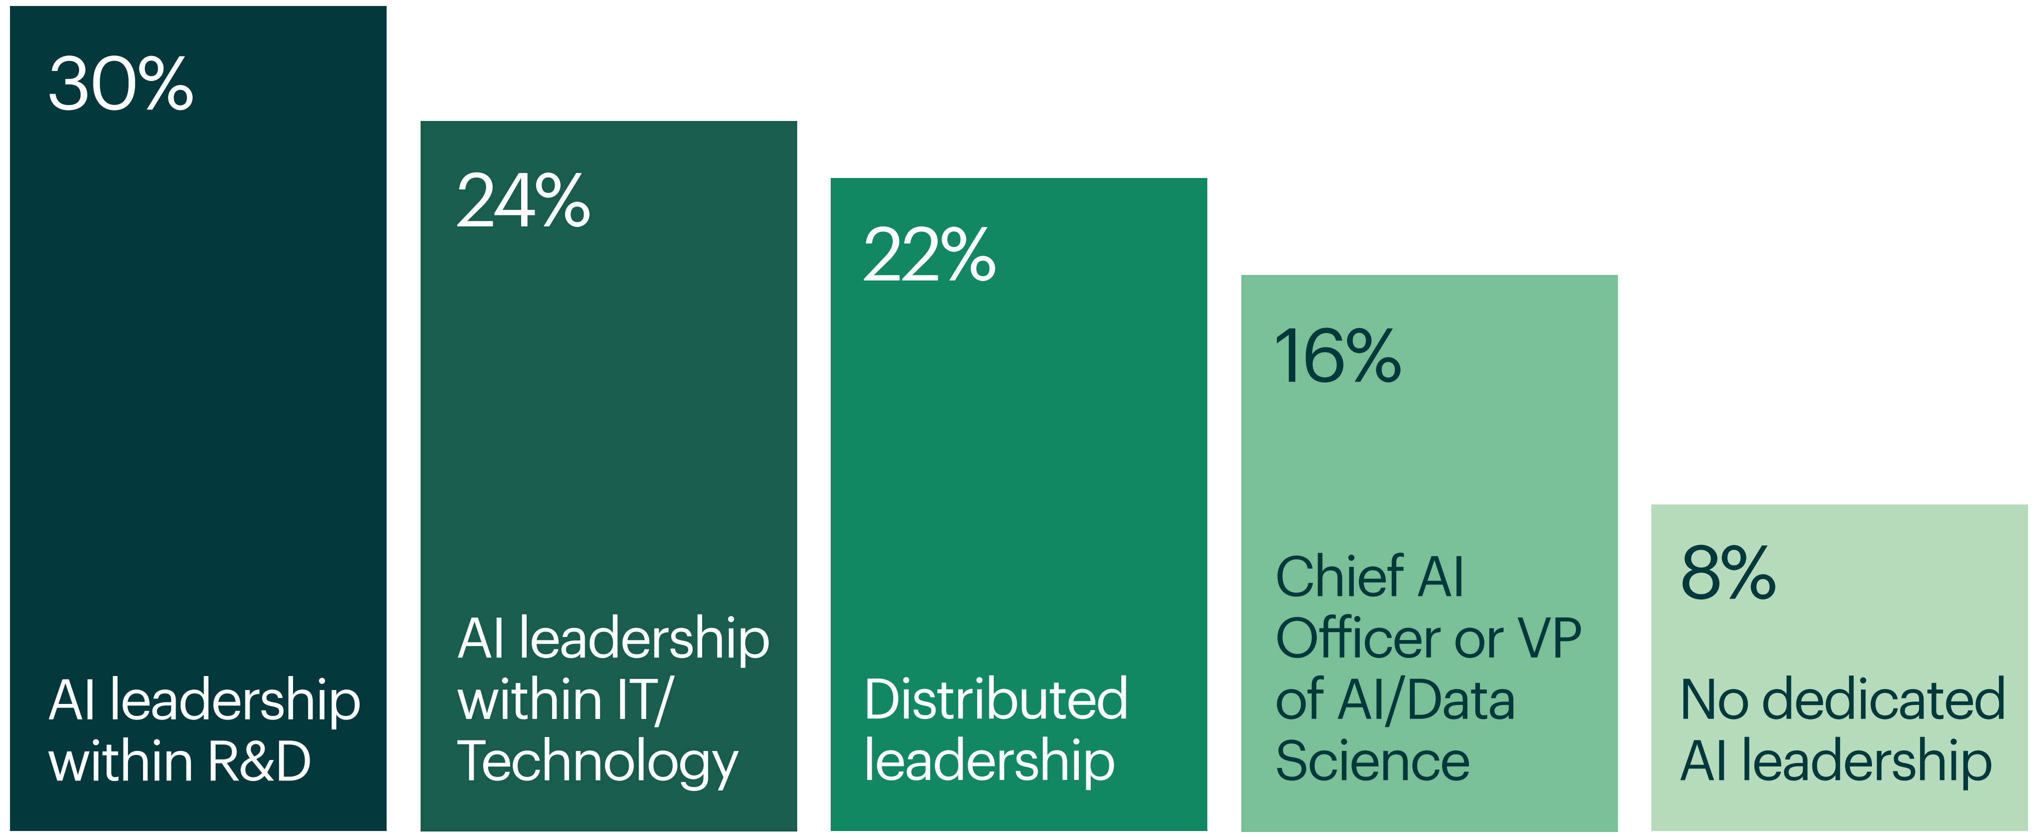


N = 104

### How AI capabilities are organized in R&D
The three most common structures

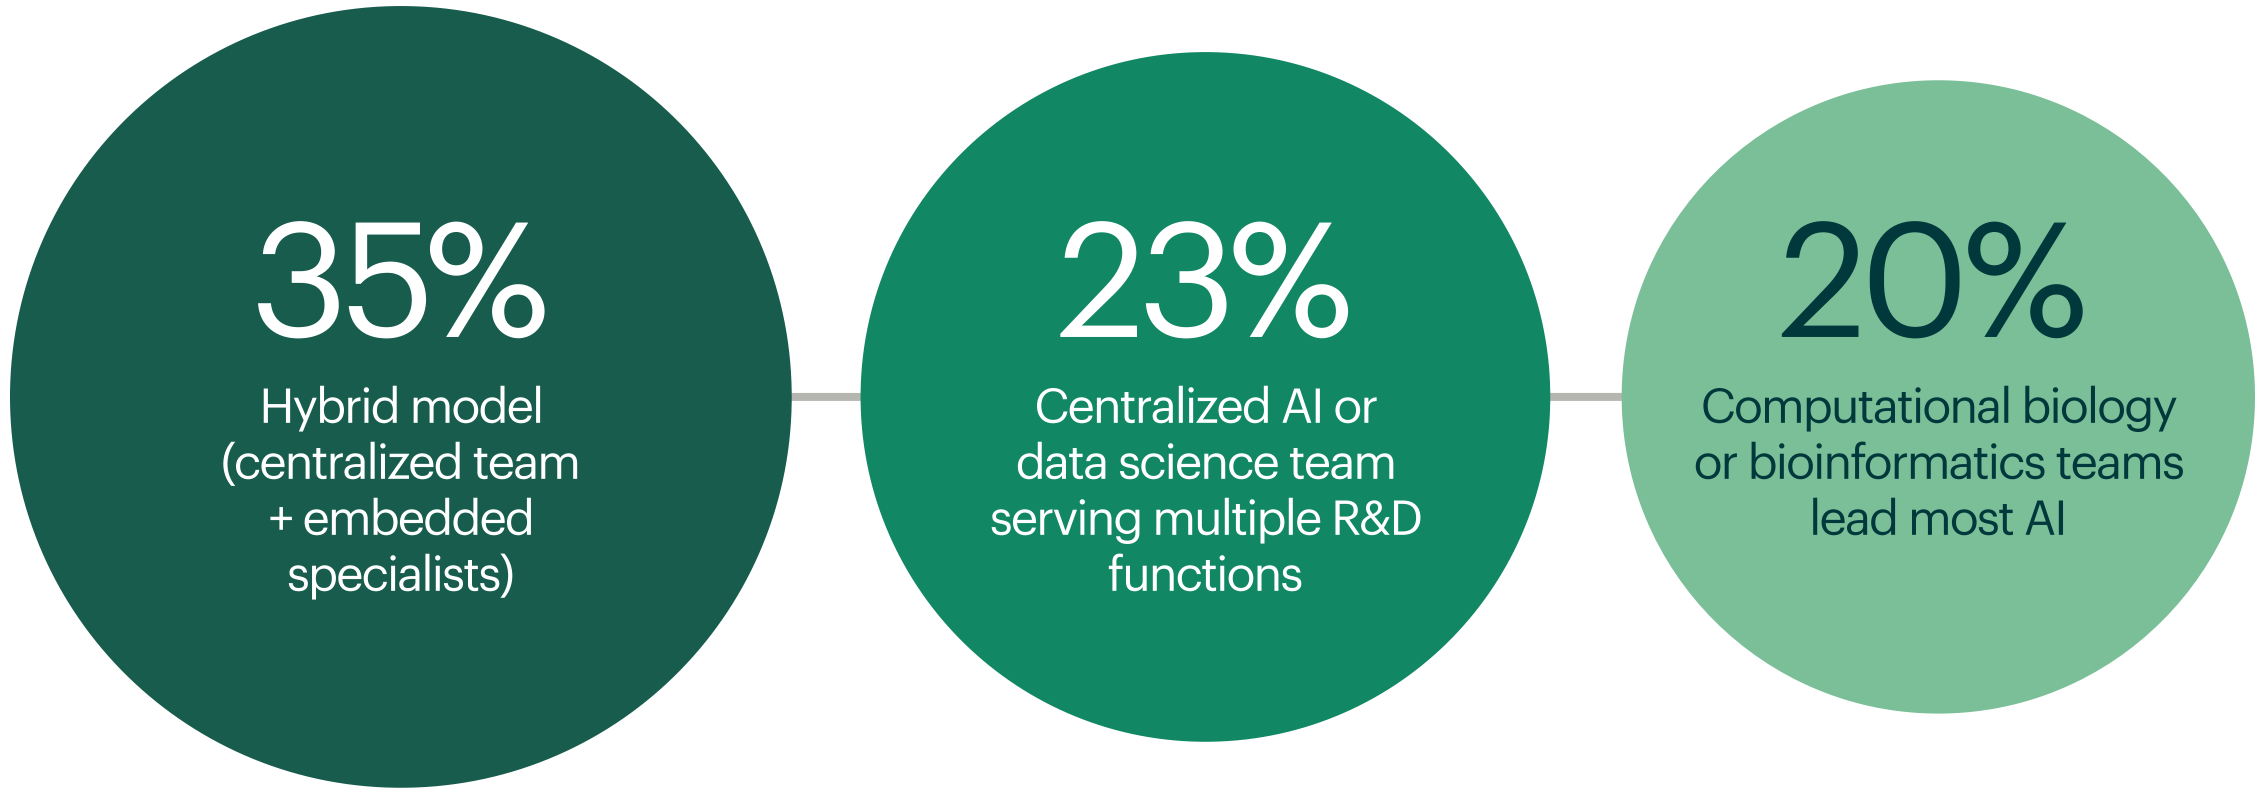


N = 102

AI leadership in biotech is shifting closer to the science. The most common model places AI leaders inside R&D (30% citing), where they can stay tied to experimental context and regulatory needs. And instead of choosing between centralized or embedded teams, organizations (35% citing) now use a hybrid approach: a core AI group for shared tools and standards, paired with specialists embedded in R&D.

The intent is to keep AI close to the bench, reduce handoffs that slow iteration, and make tools usable in the flow of real experiments. This structure supports the emerging “scientific translator” roles, the people who understand both the models and the biology they’re meant to inform.

# AI talent: Growing ‘scientific translators’ from within

Biotech’s most important AI trend is the rise of hybrid talent — the scientists who can work with models and engineers who understand the biology. Companies are building this talent internally because it’s the fastest path to useful AI adoption. Internal upskilling (67% citing) far outpaces hiring from tech (21%), reflecting the need for people who already understand the scientific and regulatory environment.

In a complex, highly regulated environment, these scientific translators help teams connect models, data, and experiments, reducing false starts and speeding validation. They also make AI more accessible: organizations report that upskilled scientists applying AI in their day-to-day workflows create more impact than small, isolated expert teams focused on advanced architectures.

**Primary sources for AI talent**

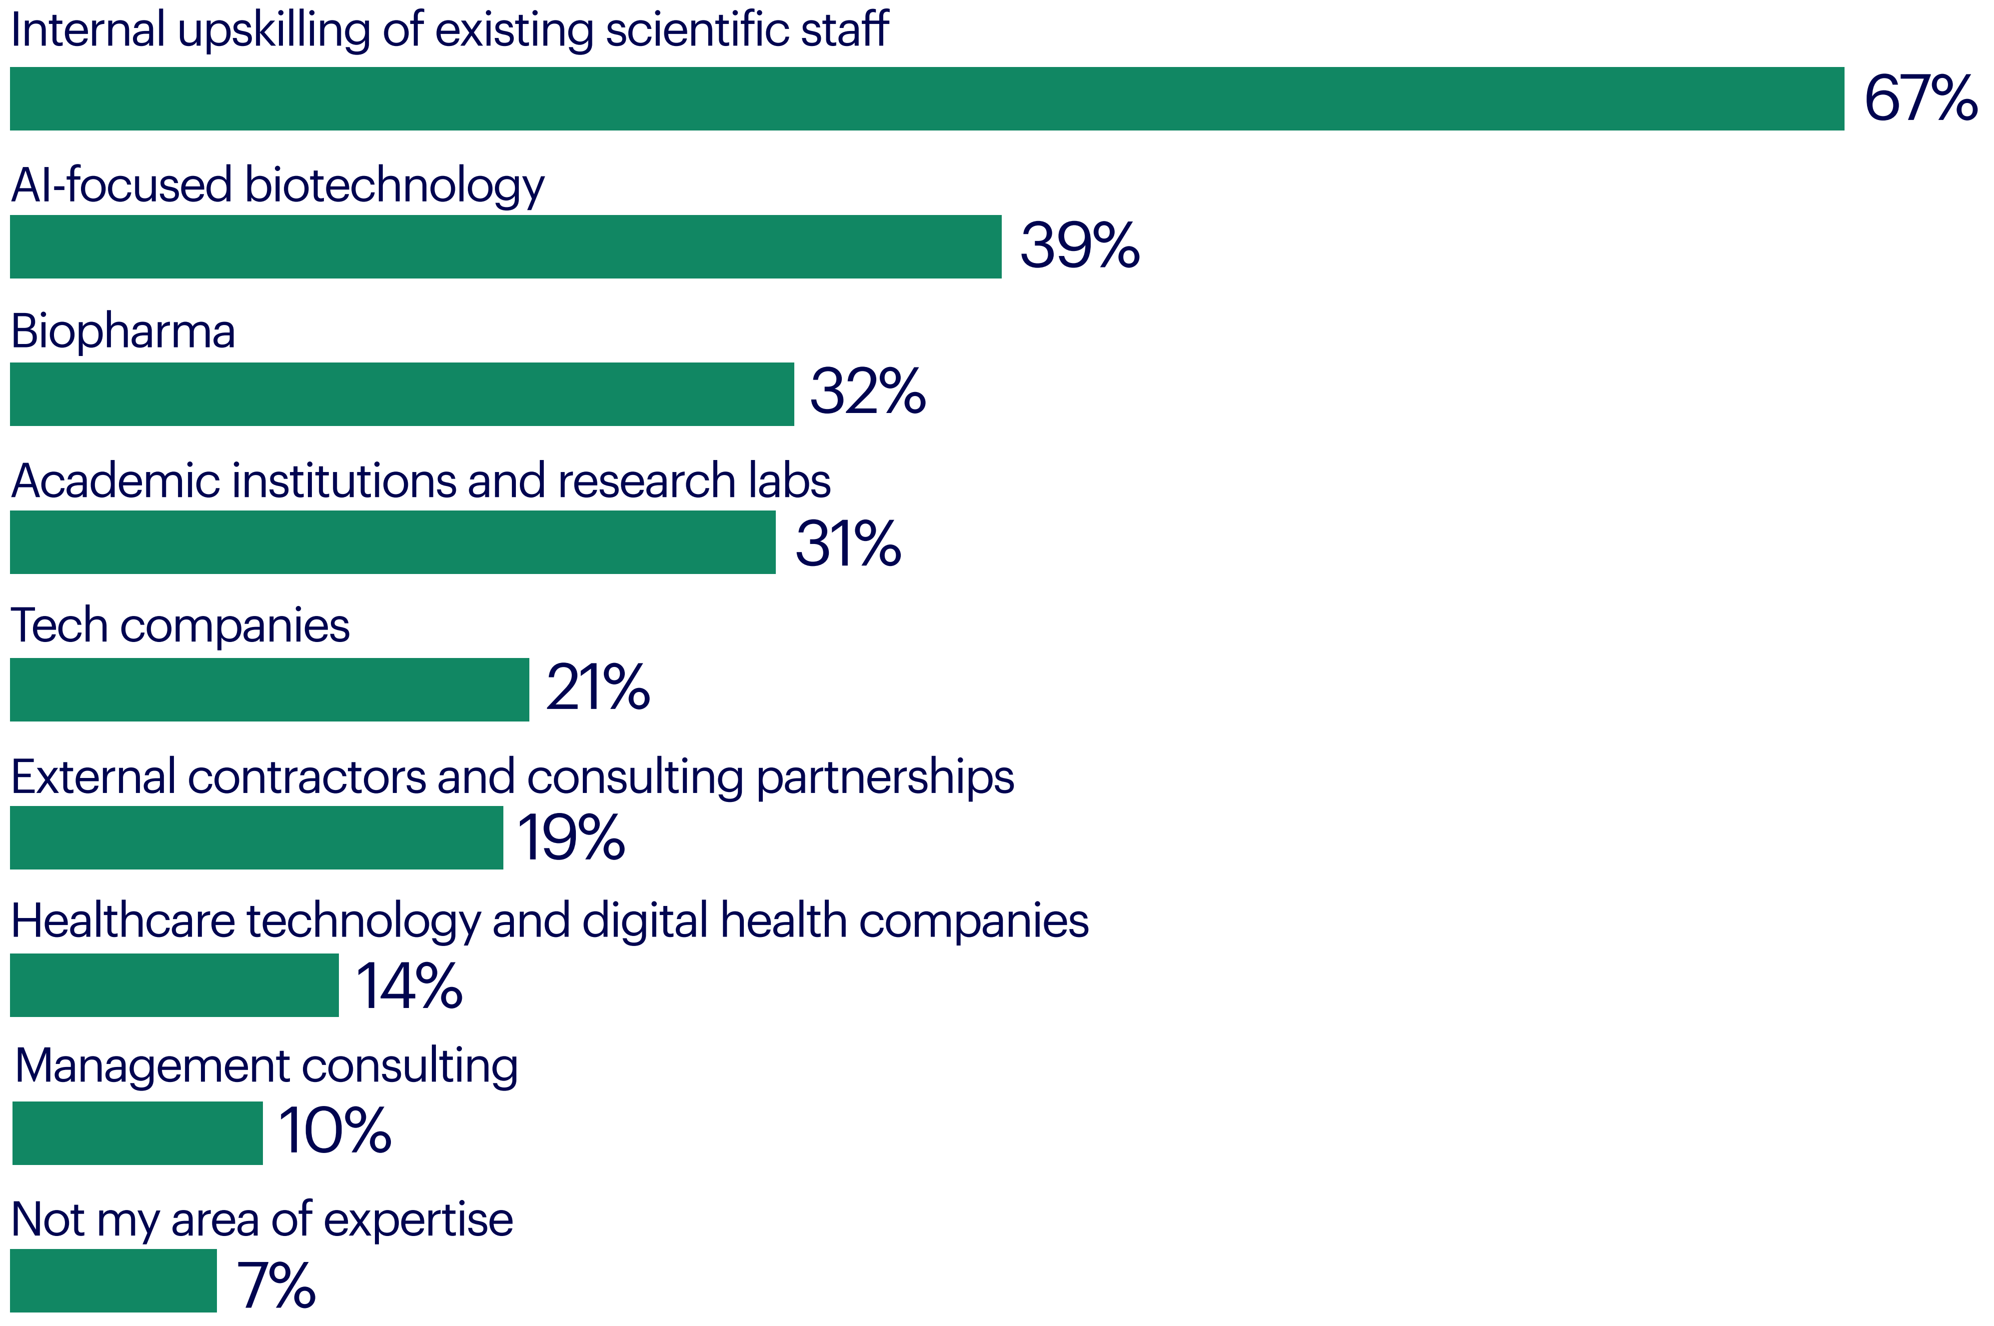


N = 104

> “We’re at a moment where computation is unquestionably here to stay in drug discovery, and the responsibility now is to balance excitement with rigorous validation. Every company is using AI in some way; what matters is how well they connect those capabilities across the organization. The biggest challenge is talent — people who can navigate both science and ML — but the progress we’re making as an industry is remarkable.

Karen Akinsanya
President, Head of Therapeutics R&D and Chief Strategy Officer, Partnerships

Schrödinger

# The roles biotech is hiring to scale AI in R&D

Data scientists with life sciences expertise and computational biologists remain top AI hiring priorities, anchoring AI efforts in real biology. But AI product managers for R&D are also emerging as a critical role as AI scales. These leaders ensure AI tools are implemented correctly, adopted by scientists, and tied to real scientific and business needs. They coordinate across external AI vendors, internal model teams, data infrastructure, IT, and R&D leadership — corralling the ecosystem so AI delivers measurable value rather than isolated capabilities.

> **Model capability will keep moving fast. The durable advantage is building workflows that can absorb better models without losing provenance, and building teams that can operate, measure, learn, and adapt as quickly as the science changes.**
>
> 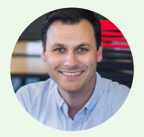
>
> Nate Gross
> Head of Health
>
> OpenAI

**AI talent priorities**

Ranked as #1 / Ranked as #2

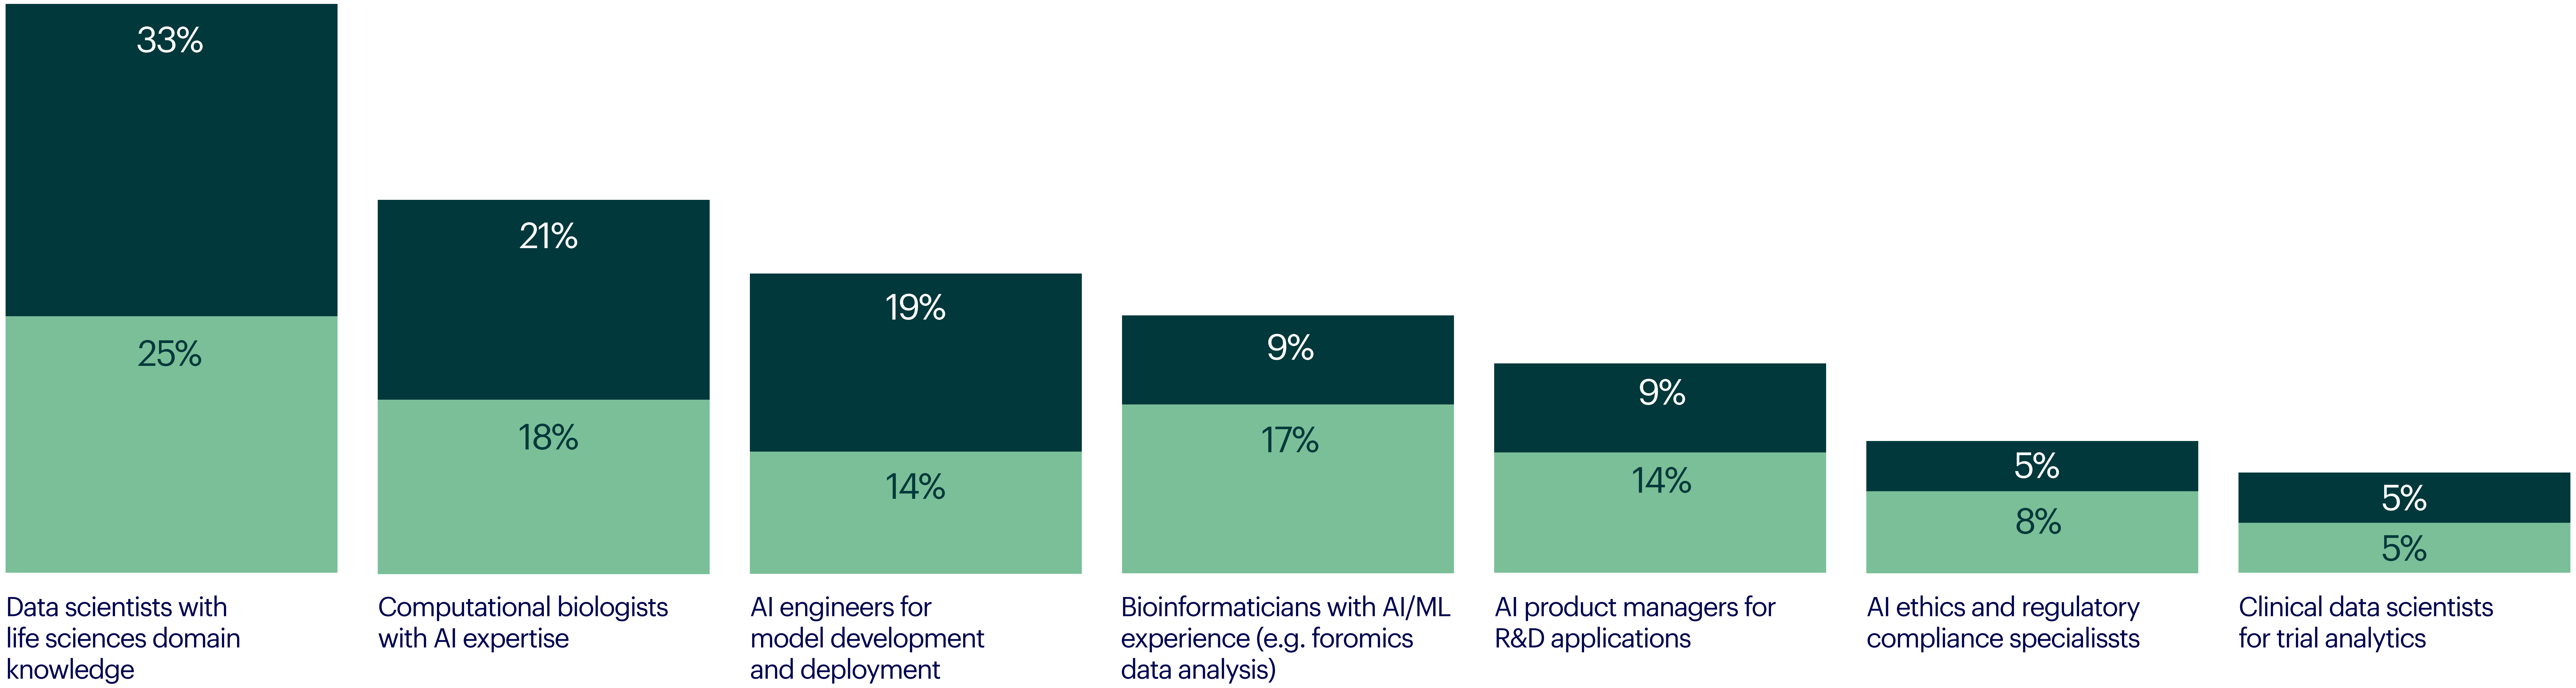


N = 104

# Operationalizing AI

## From pilots to practice, AI is becoming part of the R&D operating model

The data shows a field crossing the threshold from pilots to real operational use: scientists increasingly trust LLMs, foundation models are reaching expert-level performance, and organizations are adopting hybrid build approaches while scaling budgets to support the infrastructure beneath AI. Payment models remain fluid, signaling a market still finding its footing, but directionally, the biotech industry is kicking off a new operating model for AI.

# Scientists crossed the trust but verify threshold with LLMs

55% express moderate trust in LLM outputs, while 66% report an increase in confidence over the past year.

### Trust levels for LLM outputs

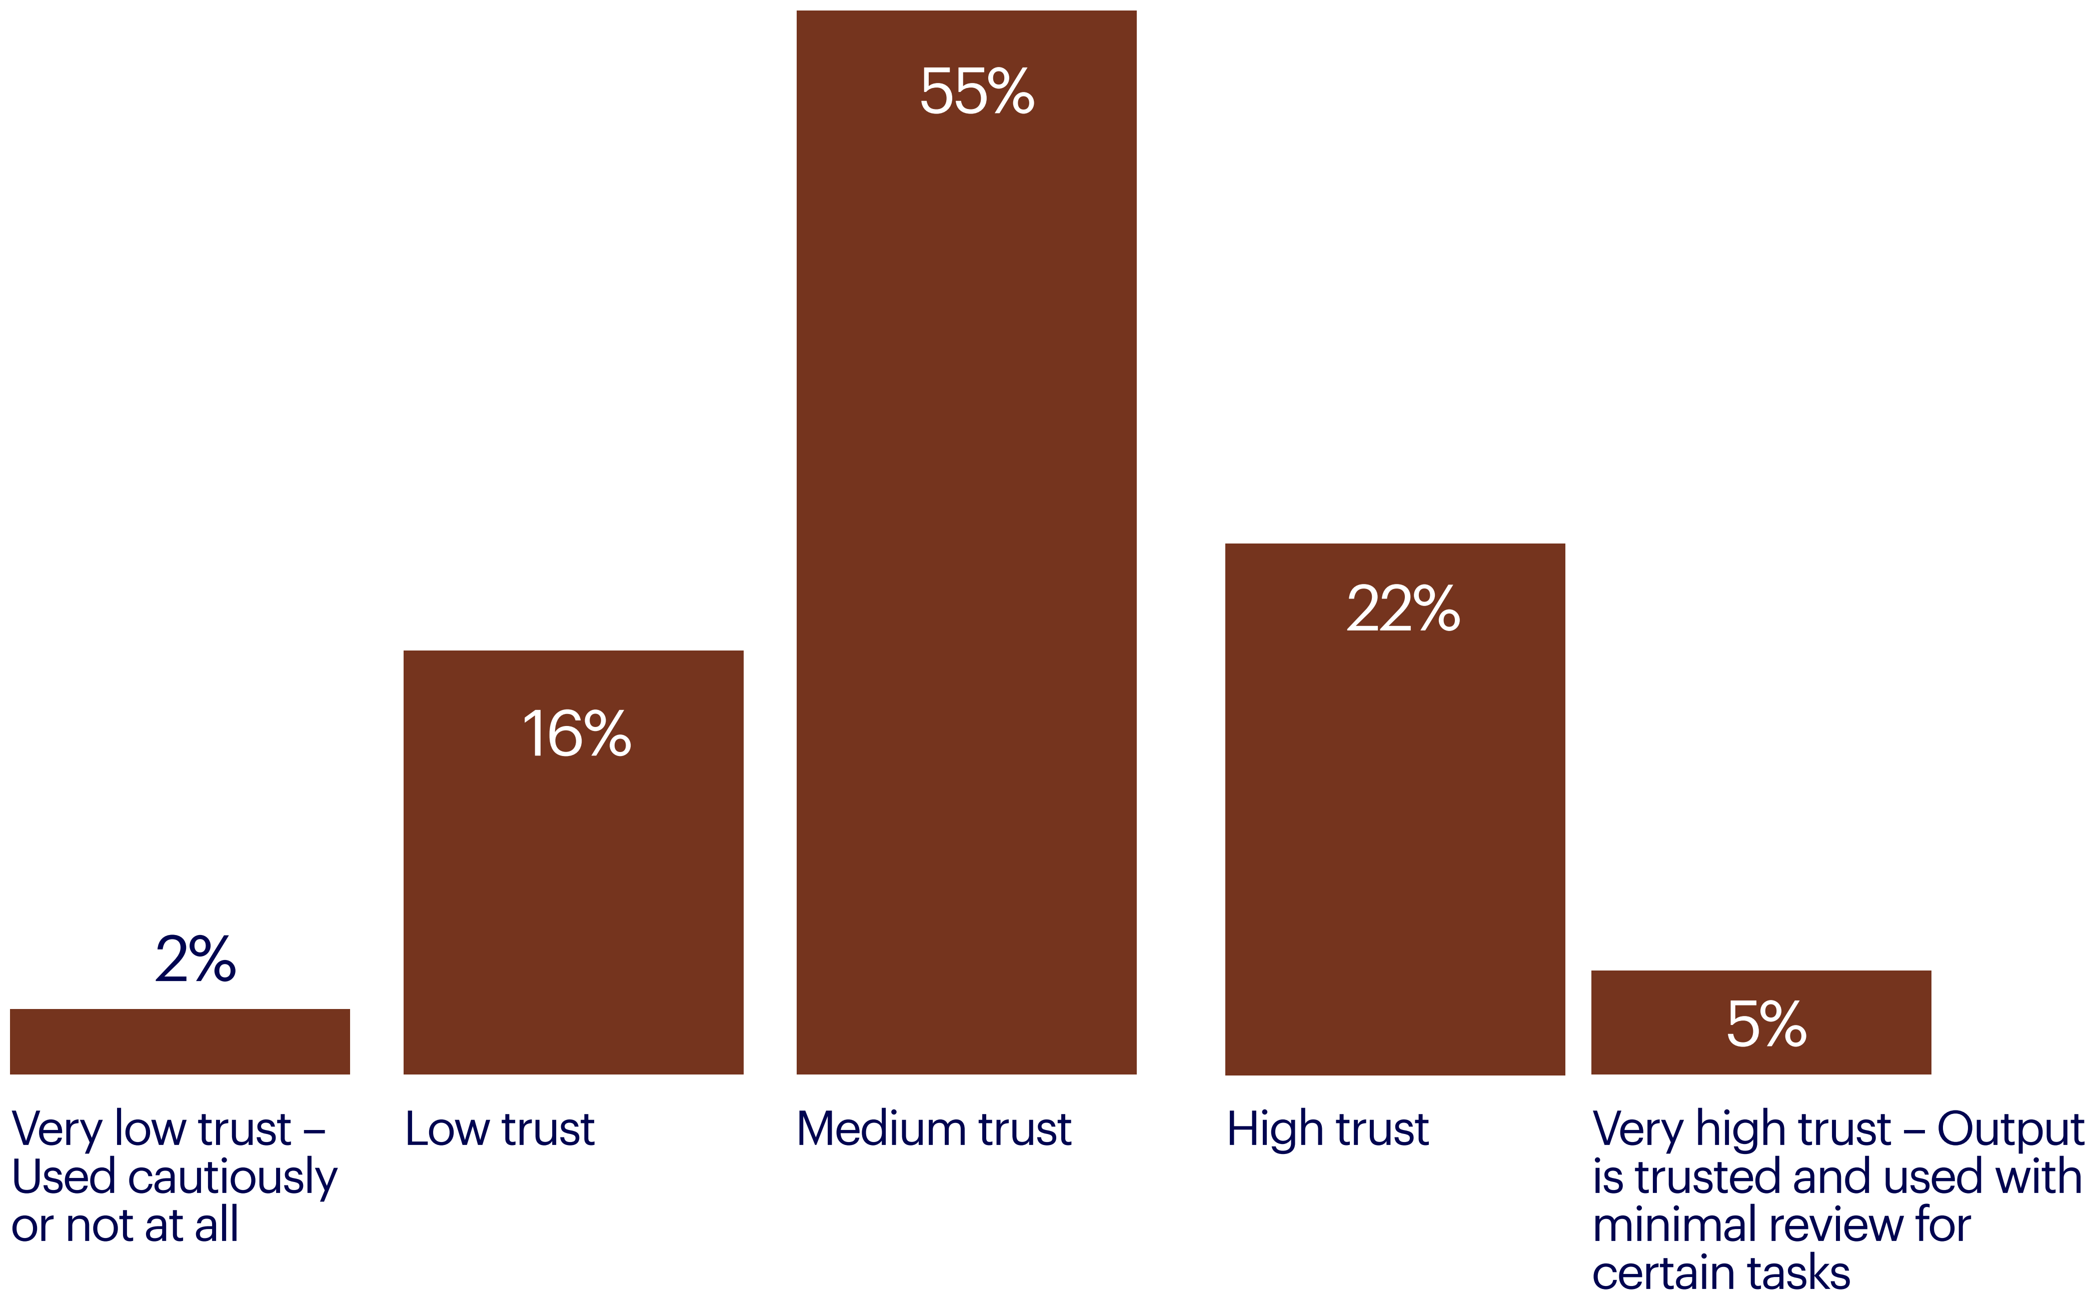


N = 104

### YoY changes in trust of LLMs' output

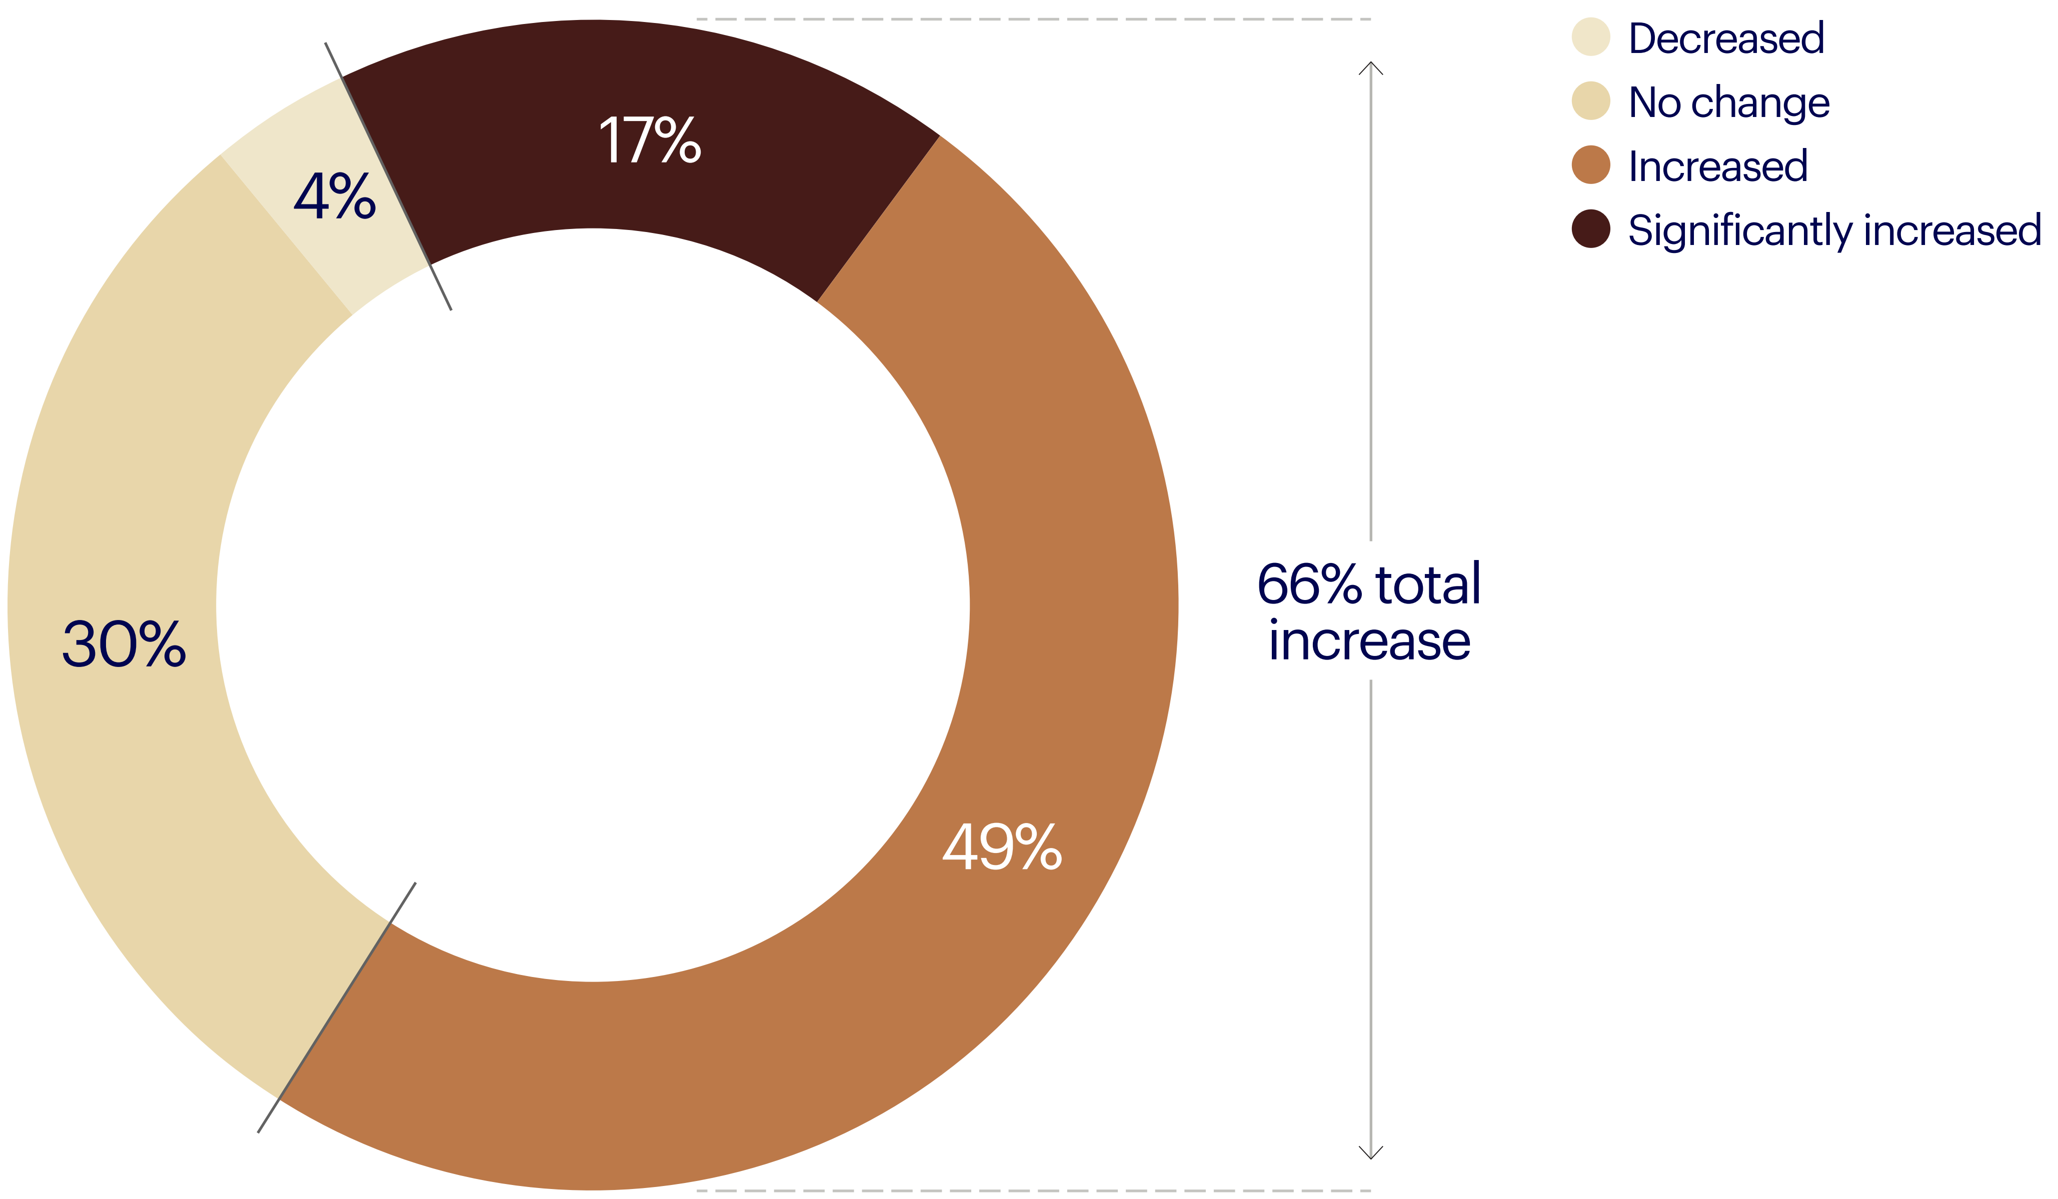


N = 104

The report's findings on trust — moderate overall, rising year-on-year — match what we see in practice. The most effective teams are those that keep humans firmly in the loop while letting AI stretch the design space. As trust moves into a "trust but verify" zone, the biggest gains come from setups where AI supports scientific judgment, not replaces it. Scientists still set the goals, watch for issues like spurious correlations or batch effects, and use their domain expertise to decide when an AI output is reliable and when it's simply a hypothesis worth testing.

# OpenAI dominates, Anthropic is surging

OpenAI has a strong wedge in biotech, used by 85% of survey respondents, with Anthropic and Mistral showing the highest growth (2.5-3x year-over-year rise). With Gemini's recent launch, along with its strong scientific benchmarks, it will likely see a rise in adoption this year.

**Generative AI vendor selection**

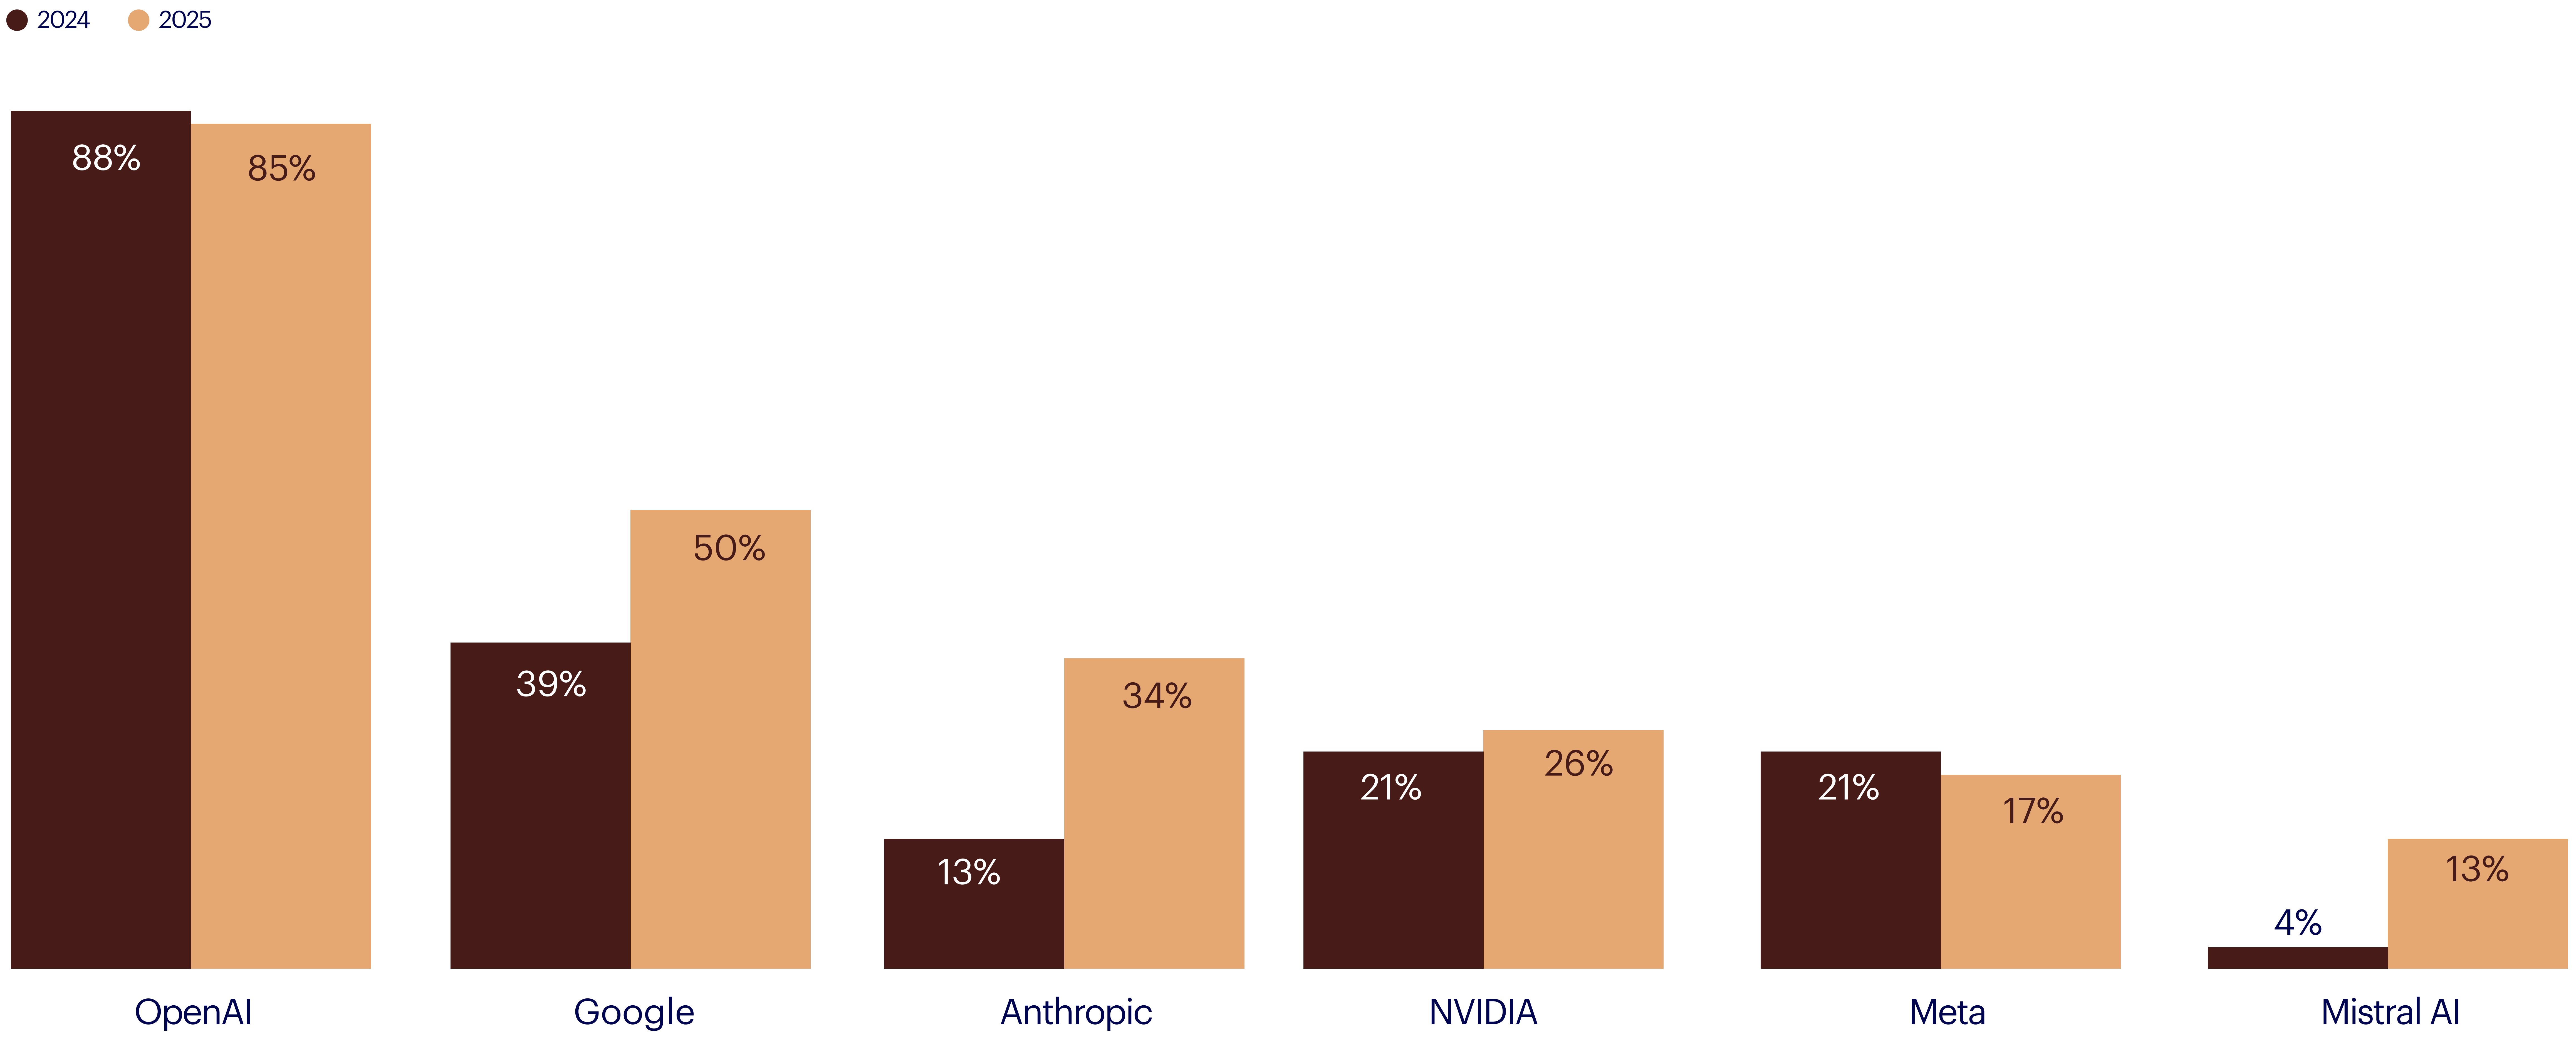

# Build vs buy vs fine-tune: Biotech adopts a hybrid builder's mentality

No single approach — building everything, buying everything, or relying solely on open-source — can keep up with the pace and complexity of AI in R&D. The field is settling into a hybrid model: 60% buy commercial AI applications, ~55% build or fine-tune models in-house where their biology is unique, and many adapt open-source tools as a starting point.

Teams need to move fast without reinventing the wheel, preserve scientific rigor in regulated environments, and build only where proprietary biology creates differentiation. It also signals early movement toward more open, collaborative norms, with organizations increasingly sharing model components, standards, and tooling as they build toward AI-native R&D systems.

### Sourcing & developing AI apps for R&D

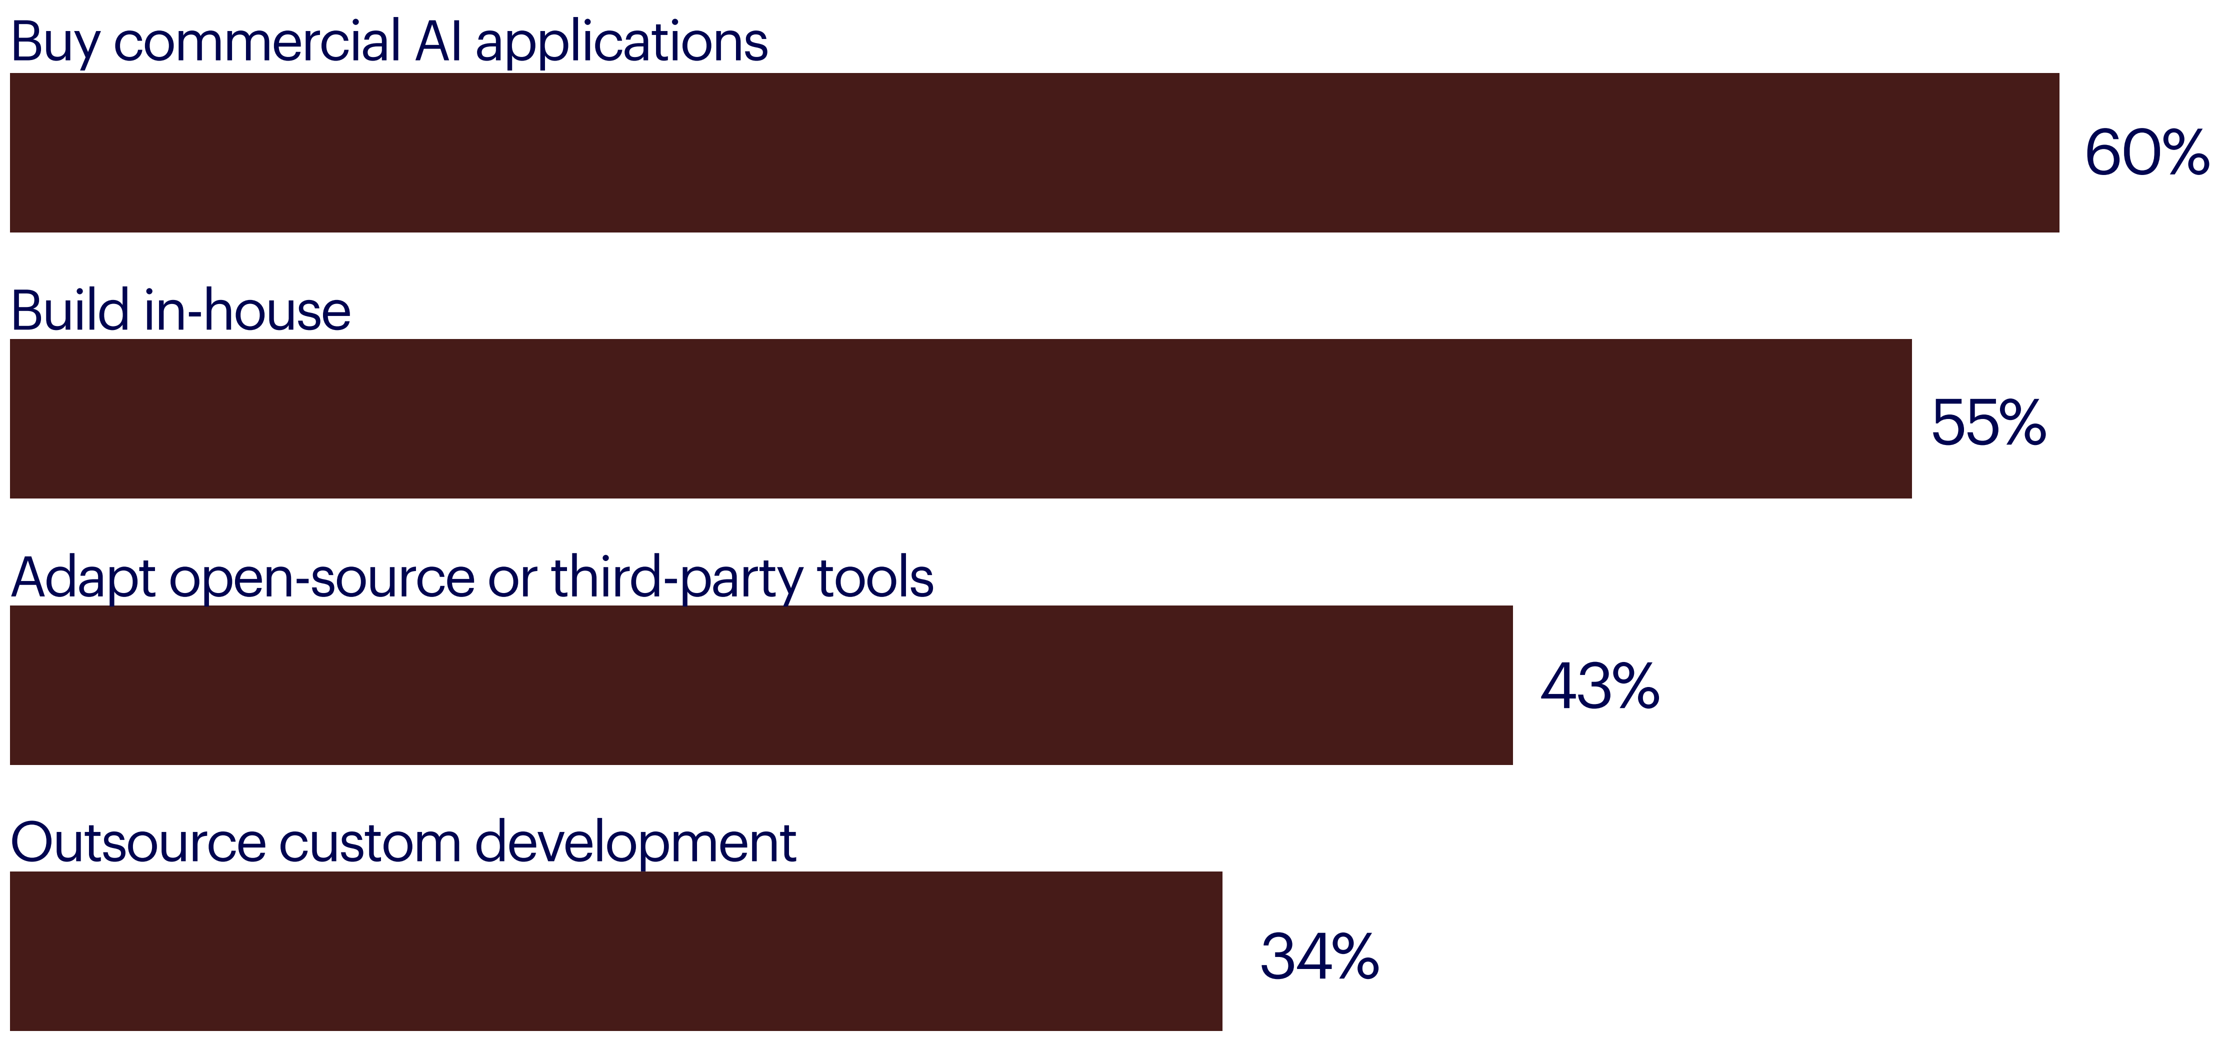


N = 102

### Sourcing & developing scientific models for R&D

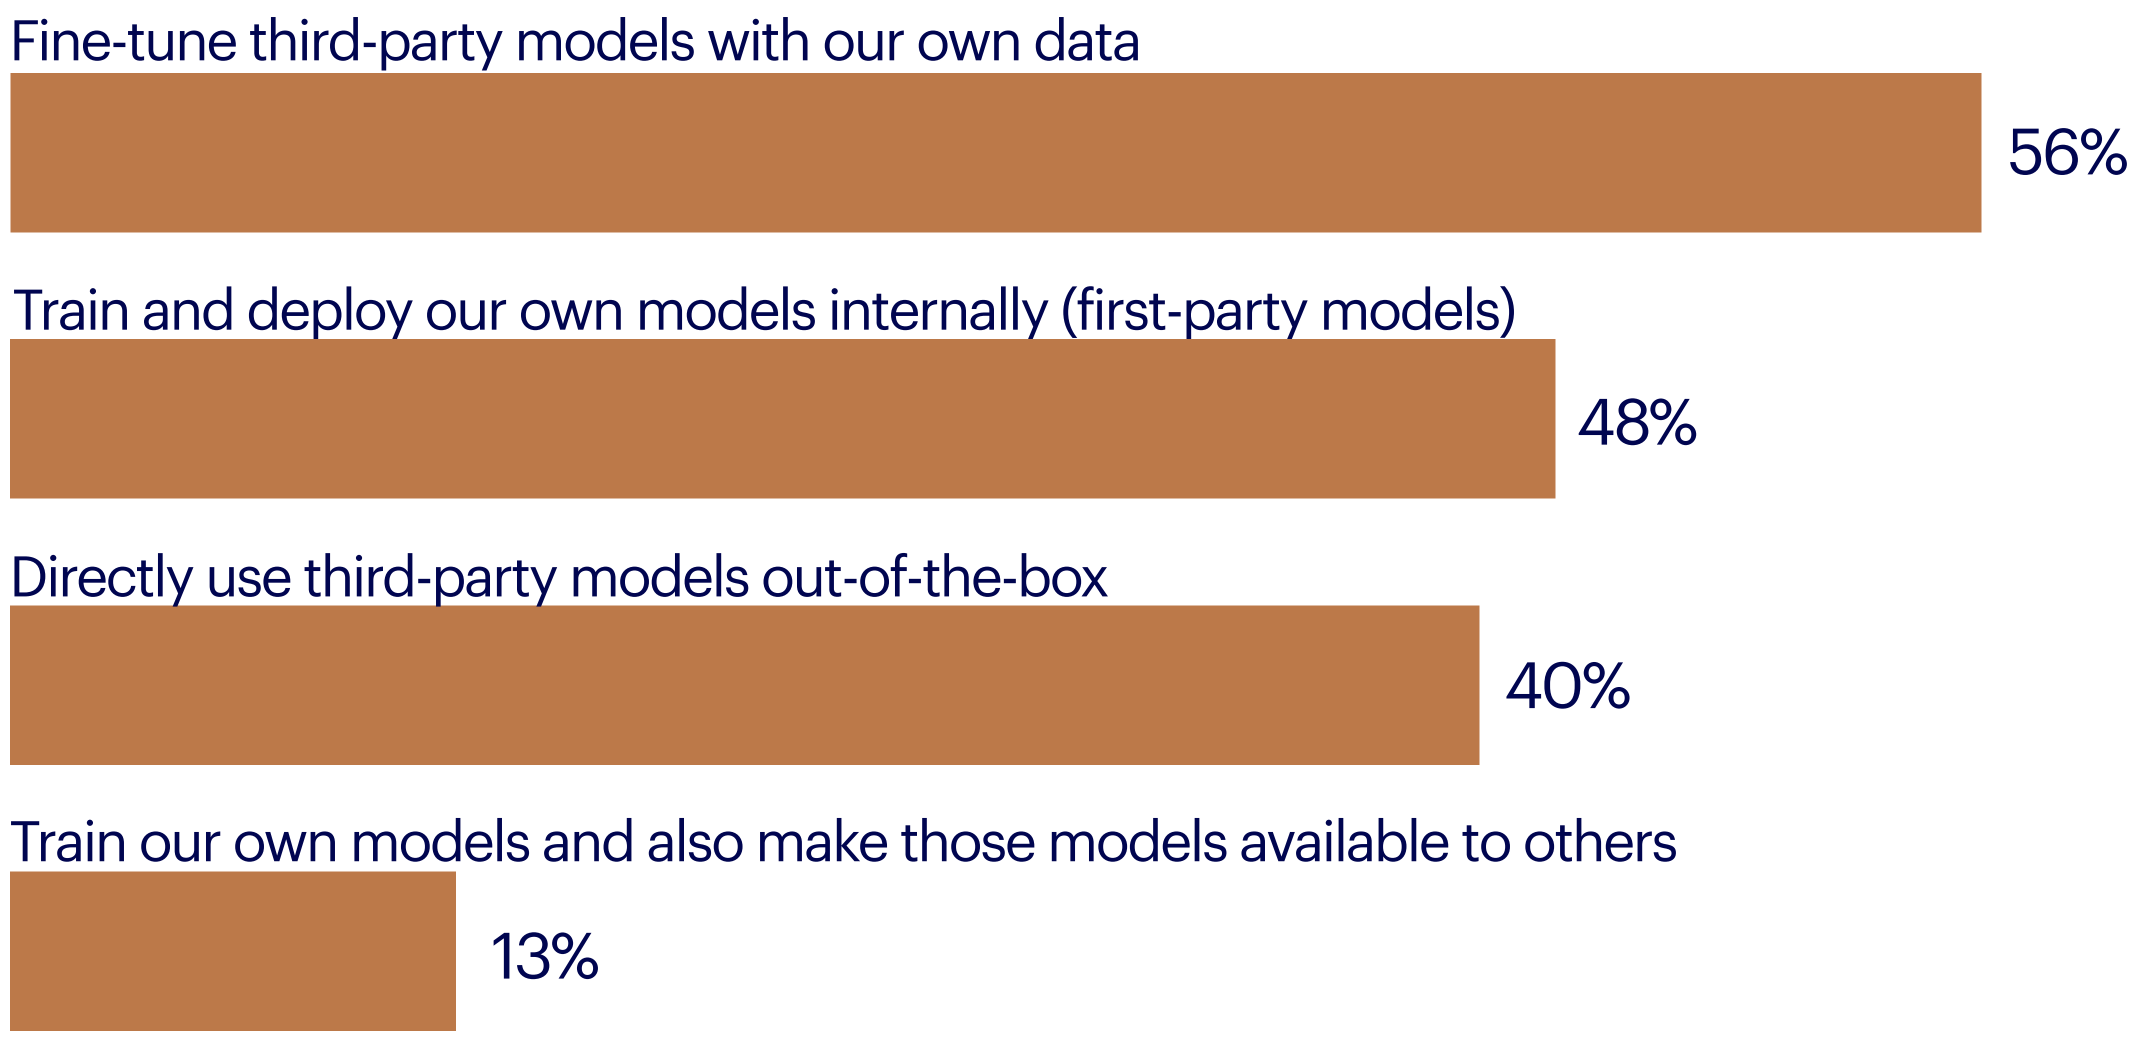


N = 103

# Budgets show AI has moved from pilot to platform, growth is accelerating

## Percent of R&D tech budget allocated to AI

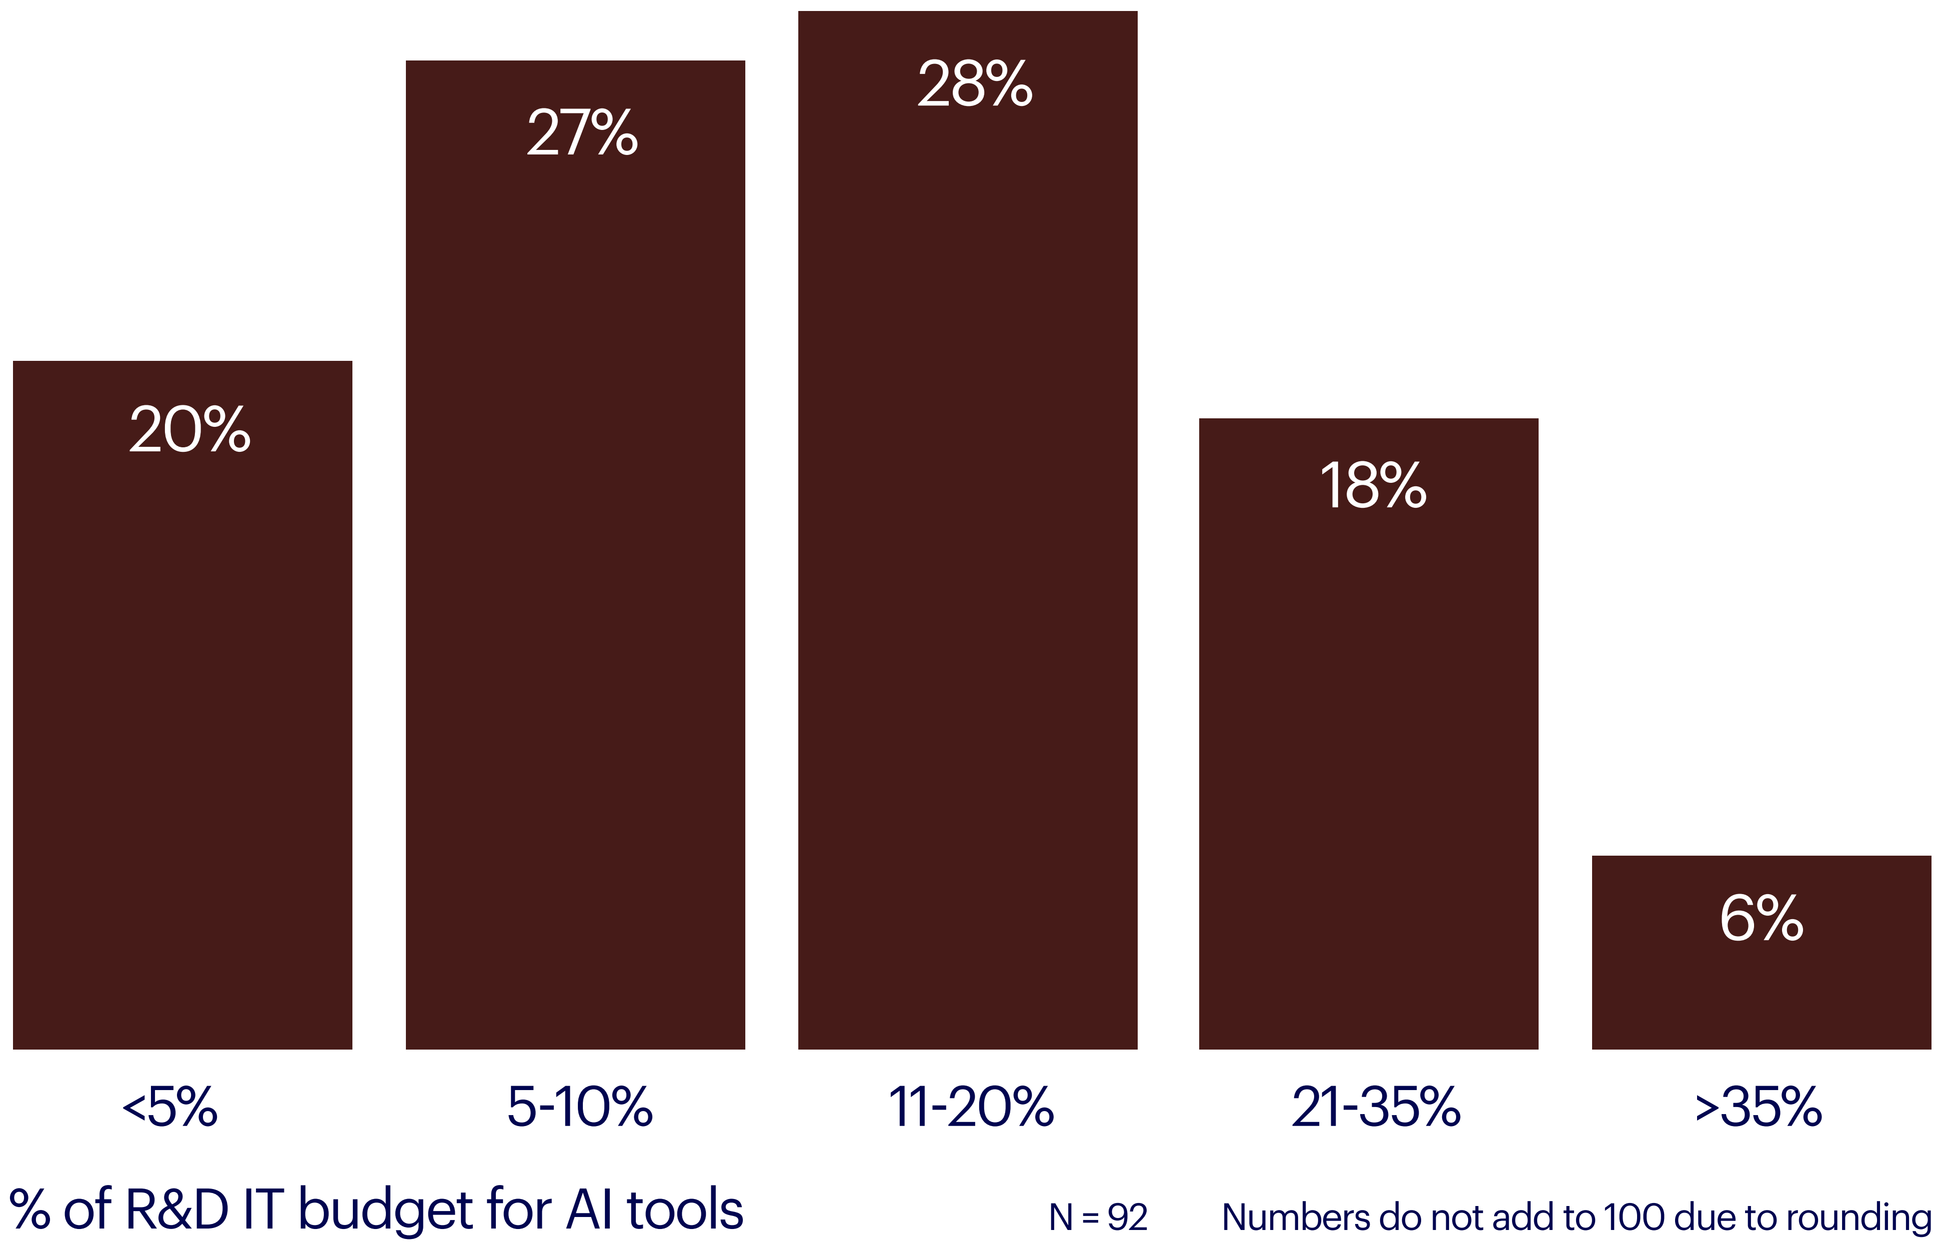


N = 92 Numbers do not add to 100 due to rounding

## YoY change in AI budget

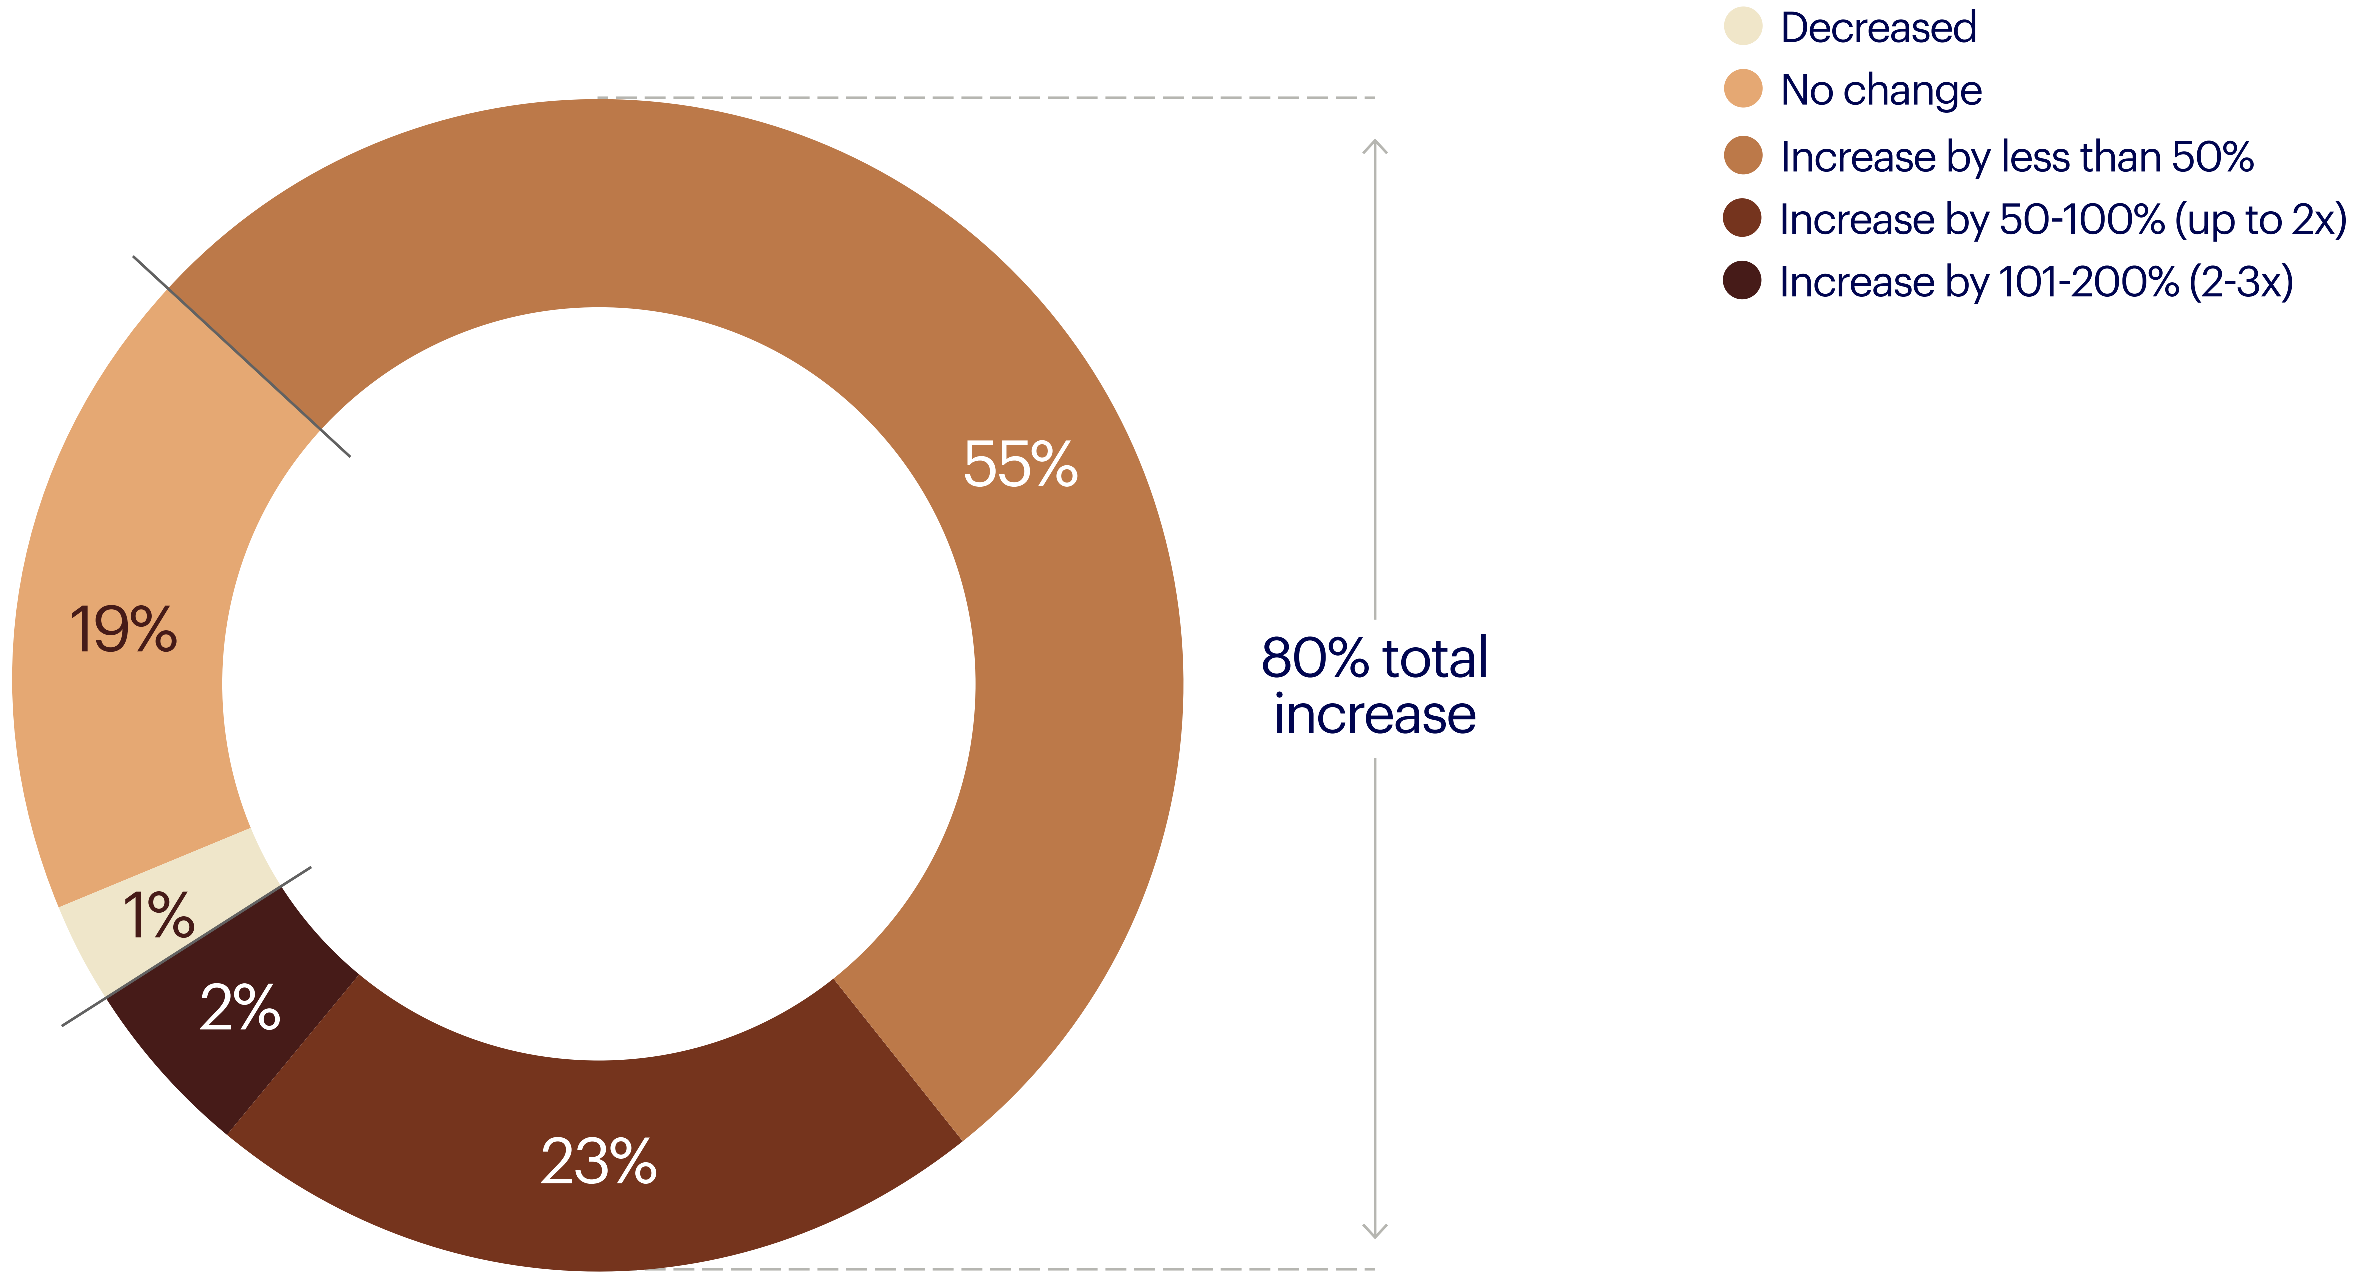


N = 97

Most organizations are scaling beyond early proofs of concept and allocating meaningful portions of R&D budgets to build data infrastructure, integrate copilots, and expand scientific modeling capabilities.

AI investment in biotech is accelerating sharply. 80% of organizations plan to increase their AI budgets in the next 12 months, with 23% expecting to double or more. Only 1% anticipate a reduction.

# Pricing models for AI are diverse, with no dominant model

There's no single dominant pricing model, showing how fluid the AI market still is. Usage-based billing — whether token-based (26% citing) or compute-based (22%) — is becoming mainstream, especially for LLM-driven tools and APIs. Only 9% organizations report value-based contracts, suggesting ROI alignment is still aspirational, not yet standard.

**How companies pay for AI applications**

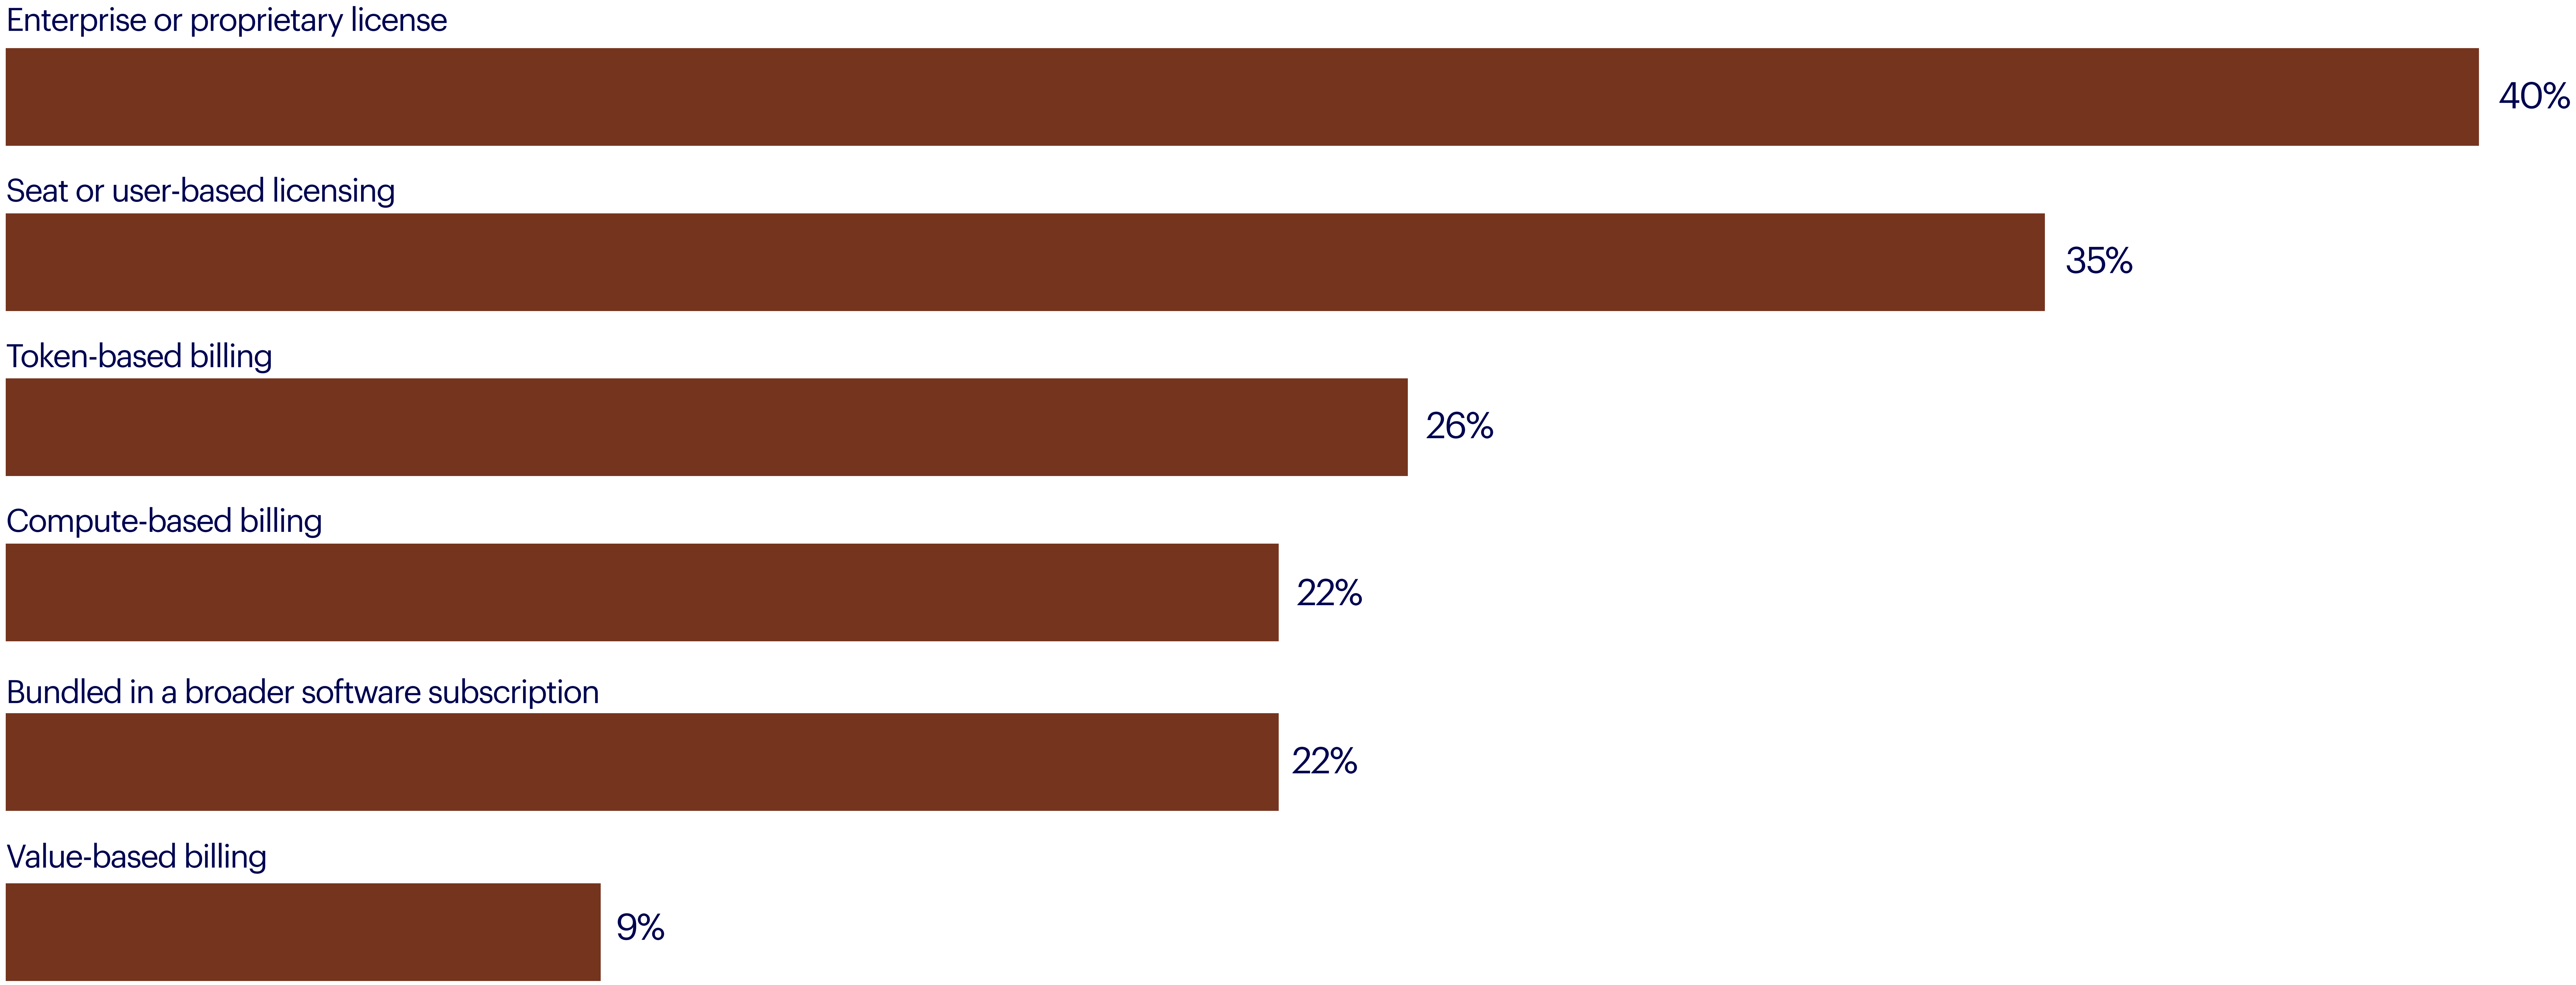


N = 104

Benchling's mission is to unlock the power of biotechnology. Founded in 2012, Benchling provides a unified, cloud-based platform trusted by more than 1,300 biotech companies worldwide, from pioneering startups to global biopharmas like Merck, Moderna, and Sanofi. Benchling's products help scientists capture, connect, and analyze data across the R&D lifecycle. With Benchling AI, scientists use agents and models directly in their workflows, connected to structured data. The result: faster teams, better molecules, and breakthroughs for all.

Trusted by the world's leading biotech companies

sanofi moderna GILEAD novonesis absci

Benchling

benchling.com