ESCP
BUSINESS SCHOOL

IT ALL STARTS HERE

BERLIN | LONDON | MADRID | PARIS | TURIN | WARSAW

© TechNova Graphics / AdobeStock

LIGHTS - Innovation/Entrepreneurship

One patient, no market: Financing N-of-1 therapies and the limits of the biopharmaceutical innovation model

ESCP Impact Paper No.2026-14-ENG

Nicolas ROBY & Frédéric JALLAT
ESCP Business School

ESCP RESEARCH INSTITUTE OF MANAGEMENT (ERIM)

ESCP Impact Paper No.2026-14-ENG

One patient, no market: Financing N-of-1 therapies and the limits of the biopharmaceutical innovation model

Nicolas Roby*
Frédéric Jallat**

ESCP Business School

Abstract

It is now possible to build a drug for one person. Genomics and RNA chemistry have advanced far enough that a treatment can be designed against a single patient's mutation, manufactured in a matter of weeks, and given before the disease runs its course. *milasen* and *atipeksen* proved it in the clinic and the problem is no longer scientific. For every patient who has received such a therapy, thousands carry variants that could be corrected the same way and will not be, because no part of the financial system is built to pay for a medicine whose market is one human being. This paper argues that the gap is structural rather than temporary: venture capital, insurance, and regulation each rest on the assumption of a patient population, and a population of one satisfies none of them. We examine why the existing financing mechanisms fail on their own terms, why computation and machine learning are reshaping the cost side of the equation faster than the financing side, and why the value of an N-of-1 program may lie less in the patient it treats than in the platform it validates. We conclude on access, because a field built on heroic individual cases will concentrate its benefits among the already-advantaged unless it is designed not to.

Keywords: N-of-1 therapies, antisense oligonucleotides, rare disease financing, computational drug design, platform innovation

*Msc Pharmaceutical management & biotechnology, ESCP Business School
**Professor, ESCP Business School

One patient, no market: Financing N-of-1 therapies and the limits of the biopharmaceutical innovation model

Introduction: a scientific revolution without a business model

In 2019, a team at Boston Children's Hospital treated a six-year-old girl with a drug that existed only for her. *milasen* was an antisense oligonucleotide designed against the exact splicing defect driving her Batten disease, and it went from concept to dosing in under a year. The frequency of her seizures fell sharply. The result mattered well beyond one child, because it settled a question the field had circled for a decade: a therapy built for a single patient was no longer a thought experiment (Kim et al., 2019).

The cases that followed made the same point in other diseases. *atipeksen* was built for a child with ataxia-telangiectasia; *jacifusen* and *afnersen* for two forms of amyotrophic lateral sclerosis (ALS); *valerioasen* for a severe childhood epilepsy. Each used the same basic move, a short synthetic oligonucleotide correcting a faulty gene before it becomes a faulty protein, and each compressed a process that normally takes a decade into months (Kim et al., 2023; Bou-Jaoudeh et al., 2025).

Regulation has started to bend toward this reality. The FDA issued guidance for individualized antisense oligonucleotides (ASOs) in 2021 and went further in 2026 with its Plausible Mechanism Framework, which allows therapies aimed at different mutations in the same gene to be assessed under a single, more flexible evidentiary logic (FDA, 2021; FDA, 2026). The N=1 Collaborative (a consortium of clinicians, scientists, drug developers, industry experts, and families) has begun publishing methods, control sequences, and shared data so that each new program does not start from zero (Lauffer et al., 2024).

Financing has not moved at all. *milasen* and *atipeksen* each cost roughly USD 1.6 million to develop, and both were paid for by philanthropy rather than by any pharmaceutical company or biotech fund (Bou-Jaoudeh et al., 2025). There is still no reimbursement code for a single-patient drug, no insurance category for a therapy with no comparator arm and no Phase 3, and no pricing logic for a medicine whose entire market is one person. Medicine has learned to intervene at the level of the individual while the institutions that pay for it still think in populations. Our argument is that this mismatch is not a gap the industry will quietly close. It is the predictable output of a financing model built on scale, replication, and market size, and it will persist until that model is rebuilt around something else.

The scale logic of biopharma

Everything about drug development assumes the cost will be spread across many patients. Development is expensive and mostly comprises fixed costs, and the prospect of revenue from a large treated population is what makes the upfront risk rational. The headline figures vary by method, from a median near USD 985 million to estimates of USD 2.6 billion, but the order of magnitude is the point (DiMasi et al., 2016; Wouters et al., 2020). That single assumption shapes how venture investors price an opportunity, how payers model budget impact, and how regulators decide what counts as enough evidence.

The orphan drug market looks like an exception but is not one. Rare disease drugs are commercially attractive precisely because they still serve identifiable cohorts and can command high prices; Evaluate Pharma projects that orphan products will exceed USD

400 billion in sales and reach around a fifth of the global prescription market by 2032 (Evaluate, 2026). Vertex's cystic fibrosis franchise and Johnson & Johnson's Darzalex are "nichebusters" that clear a billion dollars a year on populations in the tens of thousands. Small, in this context, still means a market.

N-of-1 therapies fall outside that logic entirely. They address "nano-rare" conditions affecting fewer than a hundred people worldwide, sometimes a single patient. There is no volume over which to recover development cost and no cohort with which to negotiate a price. The intervention stops being a replicable product and becomes a singular act, and the commercial machinery that makes orphan drugs work has nothing to grip.

A demand that doesn't look like one

It would be a mistake to read "one patient per program" as "a problem too small to matter." The aggregate is large: more than 7,000 rare diseases have been described, but approved orphan drugs reach only about 5% of them, leaving thousands of conditions with no treatment at all (Rodriguez Carstens et al., 2026). A large share involve splice-site mutations, intronic variants, and gain-of-function errors that oligonucleotides can in principle correct and that conventional pipelines never reach. Schork and Goetz (2025) estimate that 5% of the global population carries rare variants of uncertain clinical significance, a pool that will only grow more visible as genetic sequencing becomes routine.

n-Lorem makes the latent demand concrete. The foundation, which supplies ASOs to nano-rare patients at no cost, received more than 330 applications in its first five years, approved over 160 patients, treated more than 30, and now files roughly ten investigational new drug applications (INDs) a year (n-Lorem Foundation, 2025). That is the throughput of one charity working under tight constraints. The treatable population is too fragmented to behave like a market and too large to stay invisible, and whole-genome sequencing in routine care will only widen the distance between those who can be diagnosed and those who can be treated.

Why the current system cannot finance these therapies

The shortfall is usually described as a matter of will or timing, as though more generosity or a few more years would resolve it. The real reason is colder: N-of-1 therapies break the load-bearing assumption of every mechanism that funds drug development, and they break each one in a different place.

Venture capital is a portfolio instrument. A fund expects most of its bets to fail and relies on a few large outcomes to carry the rest, which requires every potential winner to sit in front of a market big enough to return the whole fund. A drug for one patient cannot meet that test at any price; multiplying a single course of treatment by even a few million dollars does not produce a venture-scale return. Lo (2017) has shown that high-risk biomedical innovation often needs financial engineering beyond ordinary venture logic, and the single-patient case takes that argument to its limit.

Insurance fails for a different reason. Its entire function is to pool risk across a population whose frequency and outcomes are predictable enough to price. A nano-rare therapy offers no comparable cohort, no outcome history, and no benchmark against which to set a reimbursement rate. The outcome-based agreements payers have begun to use for gene therapies, such as the CMS Cell and Gene Therapy Access Model for sickle cell

disease, still presuppose a cohort large enough to define what a good outcome looks like, so they do not reach a condition that afflicts a dozen patients (Roberts et al., 2024).

Regulation imposes a quieter cost that outlasts approval itself. Even a single-patient drug that wins authorization carries the recurring expense of keeping it, with no revenue behind it. Fondazione Telethon pays EUR 128,100 a year simply to maintain the European authorization for Strimvelis, a gene therapy for a condition affecting one to five infants per million births (Vavassori et al., 2024). A charity funded by donations can absorb that as a standing tax on doing the right thing but cannot scale to thousands of programs.

Each of these mechanisms presupposes a cohort that N-of-1 cannot supply. The question worth taking seriously is whether the cohort can be manufactured after the fact, by pooling chemistry, regulatory precedent, and evidence across many single-patient programs, instead of asking each program to stand on its own. The rest of this paper is an attempt to answer it.

The cost lever: computation, machine learning, and mechanism-based regulation

The financing debate focuses almost entirely on how to raise more money. The more tractable variable sits on the other side of the ledger. If each program can be made dramatically cheaper to design, validate, and run, the financing problem changes shape rather than simply requiring a larger subsidy. *milasen* already showed the size of the prize: a twelve-month, roughly million-dollar program against a conventional process that takes a decade and costs a hundred times more (Kim et al., 2019; Bou-Jaoudeh et al., 2025).

Machine learning is what pushes the cost down further, and it does so at the most expensive step. The dominant per-patient cost in an ASO program is the wet-lab cycle: synthesizing candidate sequences, screening them for efficacy, and testing them for off-target toxicity. Platforms such as ASOptimizer, eSkip-Finder, and Creyon Bio now predict efficacy and off-target liability in silico, cutting the number of physical iterations needed to reach a viable construct (Hwang et al., 2024; Ryczek et al., 2025). The chemistry suits this well: it is modular, its readouts are quantifiable, and a growing share of safety inference can be done computationally instead of through animal studies. Sequence design is, in effect, becoming a software problem.

Regulation is moving the same way and for a related reason. The FDA's Plausible Mechanism Framework accepts a lower clinical bar when the molecular mechanism of a variant-specific drug is well characterized, reasoning that a predictable mechanism makes for predictable risk (FDA, 2026). Oligonucleotides are the natural beneficiaries: their target sequence is unambiguous, their backbone chemistry has been validated in earlier patients, and they are far more predictable than a novel small molecule. Computation supplies much of the evidence this framework leans on, which makes machine learning and mechanism-based regulation two halves of the same shift.

The consequence is less comfortable than the efficiency story suggests. As design gets cheaper, the binding constraint moves to financing and access. AI does not make the next *milasen* affordable to deploy; it makes more *milasens* designable, faster than the system can pay for or clear them. In the near term the technology widens the very gap this paper is about, by manufacturing viable candidates the financing model still cannot absorb. There is also a quieter limit on what AI can do here. Machine learning needs training data, and a nano-rare variant has almost none of its own. These models earn their value by

generalizing across a shared chemistry and backbone, not by learning a single disease, which means the returns from computation accumulate at the platform level and not the patient level. That has a direct bearing on where an investor can capture value, if at all.

From product to platform

The reflex is to treat each N-of-1 program as a cost that ends with the patient. A more useful view, gaining ground among practitioners, is that the program produces an asset that outlives the treatment. Every single-patient IND generates real data on a chemistry, a delivery system, dosing, and a path through the regulator, and that knowledge transfers to any subsequent therapy built on the same backbone. The product does not scale *per se* but the platform behind it could.

jacifusen is the cleanest illustration: it began as a single-patient treatment for a young woman with ALS and led to twelve additional IND applications from patients carrying the same mutation, until Ionis sponsored a Phase 1-3 trial enrolling up to 95 participants (Bou-Jaoudeh et al., 2025). A bespoke therapy uncovered a real development program and, with it, a market that had not been visible before.

That one-to-many trajectory is only half the case. The deeper value is that each single-patient IND validates infrastructure, a chemistry, a delivery vehicle, a safety dataset, a regulatory template, that can be redeployed to programs aimed at thousands of patients, not only to the next nano-rare variant. Rodriguez Carstens and colleagues (2026) make the analogous point that genetically defined rare conditions repeatedly serve as proving grounds for mechanisms that later go mainstream, as CAR-T and T-cell-engaging bispecifics did on their way out of haematology and into autoimmune disease. On this reading the return on an N-of-1 program lives in the approval timelines it shortens for every broad-population product that later adopts the same validated platform. n-Lorem already operates this way and is building a blueprint for others to follow. Whether that collective value can be made legible enough to attract private capital is the open commercial question.

What's being built, and what's still missing

A number of instruments have appeared in the last few years. None closes the gap alone, but together they sketch the outline of a system.

Table 1. Emerging financing instruments and their current limitations

Instrument	Status and limitation
Non-profit developers (n-Lorem, Fondazione Telethon, N=1 Collaborative)	Have proven individualized development works outside the VC model, but depend on philanthropy and cannot scale to the full addressable population.
FDA Platform Technology Designation (2024)	Does not yet cover N-of-1, but can accelerate review of products built on an already-validated platform.
Predetermined Change Control Plans for RNA therapeutics (proposed)	Still only proposed in the EU; would pre-approve a family of variant-specific modifications under one regulatory envelope.
Rare Pediatric Disease Priority Review Vouchers (reauthorized Feb. 2026)	Not designed for single-patient programs, but a transferable, saleable voucher could draw in otherwise uninterested investors.
Outcome-based agreements (CMS CGT Access Model, 2024)	Requires a comparable benchmark cohort, which nano-rare conditions cannot supply.
Decentralized science (BIO Protocol, VitaDAO)	Unproven at clinical scale; IP-NFT ownership is not yet legally settled.

Two of these are worth singling out. The Predetermined Change Control Plan proposed by Bou-Jaoudeh and colleagues (2025) would let a developer win regulatory approval once for a family of related therapies, whether different target sequences within the same gene or the same backbone applied to a new variant, instead of restarting the full process for every patient. It is the regulatory expression of the platform argument, and it is the single change that would do most to make the economics work. The second is the Priority Review Voucher. Today a developer of a rare pediatric disease treatment earns a voucher that can be sold, sometimes for more than USD 100 million, to a company wanting faster FDA review of an unrelated drug. A dedicated nano-rare tier offering vouchers for N-of-1 INDs would put a transferable price on work that currently has none, and would reach investors who have no other reason to look at this space.

Decentralized science is the most speculative option and the most structurally interesting. Platforms such as BIO Protocol and VitaDAO use tokenized intellectual property and community governance to assemble capital for projects below the venture threshold (Hamilton, 2023). The fit with N-of-1 is unusually clean: a disease-specific organization of researchers, affected families, and donors, paired with a non-profit developer, could turn one-off fundraising into a standing pipeline that funds treatments in sequence rather than one heroic case at a time. These groups have backed a fair amount of preclinical work. Whether the model survives contact with clinical-stage N-of-1 programs, and whether IP-NFT ownership can be made to hold up legally, is untested.

Who gets access

"Who pays" is the easier question. "Who gets treated first" is the one without a clean answer. A developer filing ten INDs a year against hundreds of applications is already rationing, and the criteria organizations like n-Lorem use are mostly clinical: is the mutation tractable, is the patient declining slowly enough to benefit in time. Those are defensible criteria. They are not equitable ones.

The patients who reach the pipeline tend to be the ones embedded in elite medical networks, with access to sequencing and the means to mobilize fundraising and attention. Families in regions without genetic centres often cannot get a diagnosis in time to apply at all. Computation sharpens this problem in a way worth naming, because the AI tools driving design cost down do nothing for the diagnostic bottleneck; a cheaper drug that still requires a tertiary academic centre to identify the target only widens the distance between the connected and the invisible. Bara and Bugnariu (2026) argue that access to frontier treatment already tracks existing social inequity, and that precision medicine will deepen it unless equity is engineered in from the start.

Financing architecture cannot be settled before access architecture and then bolted on afterward. Shared registries, subsidized sequencing in underserved regions, and transparent prioritization rules belong in the design of the system, because they decide what any amount of money will actually buy. A field that began in the name of personalization can end up reproducing the oldest pattern in medicine if no one decides otherwise.

Conclusion: innovation beyond scale

This is an unusual moment. We built the capacity to treat a disease that affects one person before we built any way to pay for it at scale. milasen worked. The pressing questions are

who funds the next case, who is chosen for it, and whether cost can fall enough that both questions lose their force.

The answer taking shape across the research, the regulatory drafts, and the non-profit experiments is that N-of-1 therapies do not need a new kind of money. They need the existing toolkit, venture capital, reimbursement, regulatory incentives, pharmaceutical infrastructure, reorganized around the platform instead of the product, and around a patient cohort that does not yet exist instead of one that does. That has to be matched by sustained pressure on cost, using computational design, mechanism-based regulation, and platform amortization to bring the price down rather than chasing revenue large enough to justify it.

Concretely, this requires regulators to codify platform-level approval for families of related therapies, payers to write single-patient outcome agreements that do not depend on a comparable cohort, and the charities that have carried this field, Fondazione Telethon, n-Lorem, the N=1 Collaborative, to be treated as proof of concept rather than as permanent infrastructure. It requires an explicit position on access, so that a therapy priced at a million dollars for the first patient is not still priced at a million for the thousandth.

milasen needed a scientist willing to try something with no precedent, a family able to find that scientist, and a regulator willing to clear the path. The next ten thousand cases need something far less heroic: a financial architecture that refuses to treat a market of one as too small to bother with, and a computational and regulatory stack that makes each market of one cheaper than the last.

References

Bara, M. A., & Bugnariu, A. I. (2026). Bioethics of equity: Inequities and social justice in access to medical treatment. *Pangeea*.

Bou-Jaoudeh, M., Piaton-Breda, G., Pereme, F., & Gilbert, S. (2025). Toward an extensible regulatory framework for N-of-1 to N-of-few personalized RNA therapy design. *Therapeutic Innovation & Regulatory Science, 59*, 505–518.

DiMasi, J. A., Grabowski, H. G., & Hansen, R. W. (2016). Innovation in the pharmaceutical industry: New estimates of R&D costs. *Journal of Health Economics, 47*, 20–33.

Evaluate. (2026). *2026 orphan drug report: A safe(r) passage amid stormy waters*. Evaluate Norstella.

Hamilton, B. (2023, July 27). Unlocking scientific innovation through decentralised science. *Stanford Law School Law and Biosciences Blog*.

Hwang, G., Kwon, M., Seo, D., et al. (2024). ASOptimizer: Optimizing antisense oligonucleotides through deep learning for IDO1 gene regulation. *Molecular Therapy - Nucleic Acids, 35*, 102186.

Kim, J., Hu, C., Moufawad El Achkar, C., et al. (2019). Patient-customized oligonucleotide therapy for a rare genetic disease. *New England Journal of Medicine, 381*, 1644–1652.

Kim, J., Woo, S., de Gusmao, C. M., et al. (2023). A framework for individualized splice-switching oligonucleotide therapy. *Nature, 619*, 828–836.

Lauffer, M. C., van Roon-Mom, W., & Aartsma-Rus, A. (N=1 Collaborative). (2024). Possibilities and limitations of antisense oligonucleotide therapies for the treatment of monogenic disorders. *Communications Medicine, 4*, 6.

Lo, A. W. (2017). Adaptive healthcare financing. *Journal of Investment Management, 15*(2), 1–16.

n-Lorem Foundation. (2025). *Five-year progress update: Nano-rare ASO programs*. n-Lorem Foundation.

Roberts, M. H., et al. (2024). Outcome-based reimbursement and the CMS Cell and Gene Therapy Access Model. *Health Affairs, 43*(8).

Rodriguez Carstens, P., Moriyama, H., & Yokota, T. (2026). From genomic diagnosis to personalized RNA medicine: Advances in next-generation sequencing and N-of-1 antisense oligonucleotide therapies for rare genetic diseases. *Genes, 17*, 318.

Ryczek, N., Łyś, A., & Makałowska, I. (2025). Integrating machine learning-based approaches into the design of ASO therapies. *Genes, 16*(2), 185.

Schork, N. J., & Goetz, L. H. (2025). From precision interventions to precision health. *Nature Communications, 16*, 5024.

Vavassori, S., Russell, S., Scotti, C., & Benvenuti, S. (2024). Unlocking the full potential of rare disease drug development: Exploring the not-for-profit sector's contributions to drug development and access. *Frontiers in Pharmacology, 15*, 1441807.

Wouters, O. J., McKee, M., & Luyten, J. (2020). Estimated research and development investment needed to bring a new medicine to market, 2009–2018. *JAMA, 323*(9), 844–853.