CELLULAR INTELLIGENCE

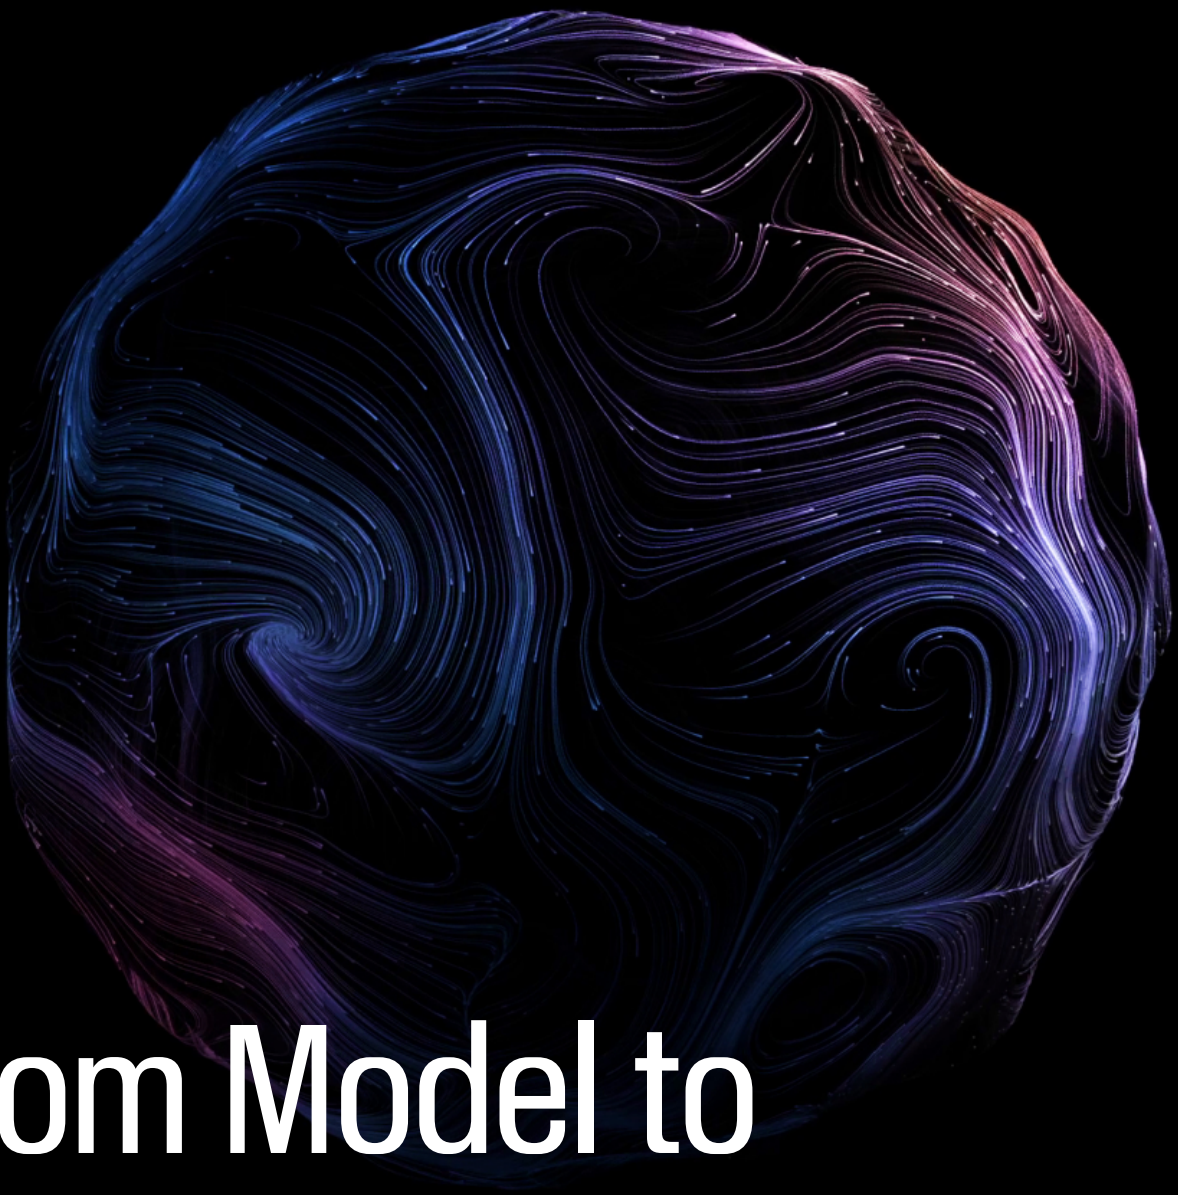

From Model to Medicine – and Back

Why we acquired a clinical-stage cell therapy, and what it means for a cell-signaling platform

JULY 2026

Disclaimer:

This paper is provided for informational purposes only. It presents scientific hypotheses, development concepts, and forward-looking plans that are subject to change as research progresses. Nothing in this paper should be interpreted as demonstrating the safety or efficacy of any product, constituting medical or regulatory advice, or implying approval by any regulatory authority. Any results, timelines, projections, and company-specific figures may be preliminary and are subject to further validation, regulatory review, and material change. There can be no assurance that the plans or objectives described will be achieved. This paper does not constitute an offer to sell or a solicitation to purchase securities.

A NOTE TO THE READER

Before turning to the scientific and business rationale, we want to acknowledge the human importance of this work. STEM-PD reflects decades of effort by scientists, clinicians, and patients. We are proud to help carry it forward, with humility and a clear sense of responsibility for its next chapter.

0.	Executive Summary
1.	Introduction to Cellular Intelligence
2.	Own the asset, control the data, close the loop
3.	Beyond STEM-PD: from one therapy to any cell therapy, and beyond
4.	Validation plan and open work
5.	References

0.

Executive Summary

Most medicines are molecules, and there is a standard playbook for producing and characterizing those molecules. Cell therapies are fundamentally different—the medicine is a bank of living cells—and so the methods for developing these therapeutics are entirely distinct.

They instead involve controlling cellular identity, often by the same mechanisms that control cells in nature[1]: by the order, dose, and sequence of the signals they see. Those are the parameters that determine what the cells become, how potent they are, and ultimately whether they help a patient [1].

Yet today, these parameters are determined largely empirically. Scientists run through many variations of a cellular differentiation protocol, measure the outcome, adjust conditions, and repeat. Cellular Intelligence was built to replace this brute-force empirical loop with a predictive one. Our AI-powered platform combines a foundation model of cell signaling, which predicts how a cell in a given state will respond to a given signal, with BioBrain, a reasoning layer that translates those predictions and other relevant knowledge into actionable insights for scientists. The goal is to predict what to change before the next experiment is run, materially reducing the number of experiments required to improve the processes of generating the desired cell type.

STEM-PD is a Phase 2-ready, Fast Track-designated cell therapy for Parkinson's disease and marks a major advance in the treatment of this disease. It also gives us the ultimate proving ground for our platform, with a concrete, high-value test: can our AI-powered platform materially reduce development time for a cell therapy? Specifically, as we transition from Phase 2 to Phase 3, we will transition from an established, manual 2D manufacturing process to a closed, scalable bioreactor process for commercial-scale supply. Such changes, while seemingly innocuous, can have major effects on cell state, potentially altering the key characteristics of the therapeutic. Thus, it is critical for us to demonstrate comparability, i.e., show that our new process generates the same therapeutic product. These experiments are costly and time-consuming, and our predictive platform will enable us to focus on the highest yield experiments, demonstrating the real-world power of our platform by potentially saving months of time and over a million dollars in costs. Acquiring STEM-PD rather than developing a clinical program from scratch allows us to make this demonstration years sooner than would have been possible otherwise while potentially accelerating the therapy's path to patients and lowering its cost of production.

Our lasting value will be the proprietary data we generate and our ability to test predictions in the wet lab against a real clinical problem.

That is one side of the flywheel: the platform sharpens the asset and potentially expedites its path to commercialization. The loop runs just as powerfully in the other direction: STEM-PD sharpens the platform. A cell-signaling model is normally trained and graded on how well its predicted cell state resembles the measured one. But resemblance is not the same as relevance. Nothing in a cell's molecular profile tells us which of its thousands of differences actually affect a clinically relevant aspect of a cell therapy. The

[1] As of December 2024, all 83 hPSC-derived products tested across the 115 regulatory-approved clinical trials catalogued by Kirkeby, Main, and Carpenter relied on directed differentiation to establish therapeutic cell identity, rather than transcription-factor-driven forward programming [Kirkeby, A., Main, H., & Carpenter, M. (2025). Pluripotent stem-cell-derived therapies in clinical trial: A 2025 update. Cell Stem Cell, 32(1), 10–37. https://doi.org/10.1016/j.stem.2024.12.005]

rich data from the STEM-PD program — both already generated and to be generated, including manufacturing process data, animal studies, and clinical outcomes — provide ground truth data for learning which parts of the cell state matter. We expect what we prove here to extend to essentially any cell therapy, all of which face similar challenges of scale, comparability, and product quality.

Owning the asset, rather than partnering on it, allows us to operate and control all aspects of this data flywheel. Ownership means we control the experiments that are performed and hold the data for both the manufacturing process and clinical outcomes. With STEM-PD in house, we set what gets measured, own every result, and act on what we learn at our own cadence and on our own terms, while maintaining our commitment to patients and the overall success of the program.

As general AI capabilities become more broadly available, our lasting value will be the proprietary data we generate and our ability to test predictions in the wet lab against a real clinical problem. Our platform supplies the general cellular intelligence, and that intelligence will be compounded by the process and outcome data we tie back in. STEM-PD is our first full-stack demonstration of that system, from model to medicine and back.

WHERE WE STAND

AREA	DEMONSTRATED TODAY	IN EXECUTION	STILL TO PROVE
PLATFORM			
Data	Large-scale sequential perturbation data.	Linked process-and-clinical data. A growing proprietary process dataset.	Calibration of which cell-state features track potency and graft performance.
Predictive modeling	Early predictive modeling.	The first model-guided process experiments.	Prospective model lift over strong baselines.
STEM-PD			
Clinical program	A Phase 2-ready Parkinson's program with prior PET evidence of graft survival.	Phase 2 on the established product.	
Manufacturing and transfer		A scalable suspension manufacturing process. A risk-based comparability package.	Transfer to other cell therapies.

1.

Introduction to Cellular Intelligence

Cellular Intelligence was founded around the premise that understanding the language of cells will enable a new era of biological design. In particular, our thesis is that the collection of high quality, high throughput perturbative data will yield machine learning models with the predictive power to steer cell fate.

In the process of development, cells differentiate from a single fertilized egg to the myriad mature cell types in the adult body. Cells do not move from the initial stem-cell state to a mature identity in a single step. During this process, cells are sequentially exposed to signals that drive cell fate specification. The effect of each signal depends on its dose, timing, and sequence, as well as the state of the cell receiving the signal. These "context-dependent" effects are the reason even a relatively limited number of signals can, via combinatorial complexity, produce all the specialized cell types in the human body [2,3].

Our platform is built to query exactly these kinds of signaling perturbations. It applies defined signals across many starting cell states, measures the resulting transitions, trains a foundation model to predict the next state, and uses BioBrain to turn those predictions and other evidence into insights and hypotheses that can be tested in subsequent experiments. The platform is therefore aimed at the natural control system of development itself: for a given initial cell state, which signal will move a cell toward (or keep it within) the desired end state.

STEM-PD is a concrete instance of the same developmental logic. A well-characterized human embryonic stem cell is guided down one branch of the developmental tree toward an A9 dopaminergic midbrain identity through a defined sequence of signals. The sections below introduce the biological principle, the platform built to interrogate it, and the clinical-stage asset in which those pieces come together; Section 2 then describes how the platform and asset can improve one another within an owned, closed learning loop.

1.1 CELLULAR SIGNALS, DEVELOPMENT, AND CELL THERAPY

In a cell therapy, the product is a population of living cells, and its therapeutic benefit depends on what those cells are when they reach the patient: their identity, purity, potency, and viability. That cellular "state" is dictated by how the cells are made.

Cells are made by a sequence of signals. A stem cell becomes a dopaminergic progenitor because it saw particular signaling molecules, at particular times, in a particular order, and its response at each step depends on the state it was already in. Change the signals, the timing, or the conditions, and you change the product. A large component of developing a cell therapy is therefore the study of how cell states change upon the application of signals — which are exactly the perturbations our platform is built to predict.

This is why cell therapies, and not just any clinical asset, benefit uniquely from our platform. A small molecule is a fixed entity defined by its chemical structure; a cell signaling model has nothing to act upon. Producing a cell therapy, however, requires defining and manipulating cell states, the very thing our model predicts, allowing our platform to improve the product itself.

1.2 OUR PLATFORM: DATA, MODEL, AND BIOBRAIN

Our platform has three key layers that enable us to predictably control cell state:

The data.
We generate our own training data at scale. Using miniaturized capsule technology, we apply a standardized panel of signals across hundreds of distinct starting cell states and read out the result by RNA sequencing. The point of this design is to capture context dependence, namely, the fact that the same signal does different things to a cell depending on the state the cell is already in. This context dependence is what has made signaling so hard to predict. Until now, there has been no systematic, large-scale experimental system to produce the data required to fully illuminate the context dependence of signaling.

The model.
The foundation model trained on our data forms a prediction engine: given a cell's current state and a signal, the model predicts the cell's next state [4]. Using modern machine-learning methods, the model predicts how cells respond to signals even in starting states it has never seen. We have built proprietary evaluations that measure the model's real-world predictive power on a tight feedback loop, so we know how well our model is able to perform and how best to improve upon it. Early tests indicate that our model is already achieving state of the art performance on real-world prediction tasks.

BioBrain.
Predicting cellular states (resemblance) is not the same as predicting what a cell will ultimately do (relevance). BioBrain is the reasoning layer that turns the model's predictions into recommendations a scientist can act on. It reasons across several inputs at once — the foundation model, a map of how cells develop, the published literature, our patent library, and our own internal and clinical data — and grades the evidence behind every recommendation. Importantly, BioBrain generates genuinely new, testable hypotheses that are unique from what can be derived from established biology; this makes it very clear when the model has made a valuable and novel contribution. That allows us to pinpoint the precise experiments with the highest potential yield to, say, improve a particular part of the cell differentiation process. BioBrain is an extensible platform that can incorporate new and varied data and knowledge streams. These include the many modalities and proprietary non-clinical data routinely generated by the STEM-PD program, as well as years of curated internal know-how, including practical parameters such as reagent cost and clinical knowledge. Building a STEM-PD-specific BioBrain is a key step towards realizing the potential of the cell signaling foundation model our platform is creating.

A large component of developing a cell therapy is the study of how cell states change upon the application of signals — which are exactly the perturbations our platform is built to predict.

1.3 WHAT STEM-PD IS

STEM-PD is an allogeneic cell therapy: A9 dopaminergic midbrain neurons grown from a well-characterized human embryonic stem-cell line and delivered into the putamen through a surgical cannula, where the transplanted progenitors mature into dopamine-producing neurons [5,6]. Its aim is to restore function by replacing the cells Parkinson's destroys [7,8,9], which distinguishes it from gene therapies or deep-brain stimulation that manage symptoms continuously but do not address the underlying pathophysiology.

Although a protocol for generating and delivering the cell therapy is established, process improvements for commercialization are required. While the current cGMP process can support Phase 2, the Phase 3 and commercial process will leverage our platform to scale up and replace raw materials with higher quality alternatives. Changes to concentrations and timing of the use of compounds, as well as substitutions for equivalent or better versions of the compounds themselves, are standard during the development of a clinical product as long as comparability to the previous product is demonstrated. The platform's job here is therefore to refine an existing, validated protocol rather than design one from scratch. This is a narrower and more tractable problem, but solving it could still have a major impact on the therapy's development.

Where it stands

STEM-PD has been tested in a single first-in-human, open-label study in eight patients. The program has since been cleared to proceed to a Phase 2 trial and granted Fast Track designation December 2025. First-in-human dosing is complete; Phase 2 has not yet begun, with the first patient expected to be dosed in early 2027 at sites in the United States and Sweden.

The Phase 1 results, now published in Nature Medicine [10], showed that the cell product and surgical procedure were well tolerated, with early, dose-related but submaximal signals of graft survival and motor improvement [11]. The cells and surgery themselves have had a favorable safety profile; the principal managed risk is the immunosuppression regimen rather than the product or the procedure. Immunosuppression was monitored throughout, and one immunosuppression-related fatal infection occurred during the study.

The Phase 2 trial tests two ascending doses, both above the prior first-in-human high dose — a response to those submaximal signals — with the second cohort gated on an independent safety review of the first; specific dose levels are held internally. The practical upper bound on delivered dose is not established: there is likely room within current surgical approaches, and emerging delivery devices could raise it further. The primary endpoint is safety; supportive endpoints are motor function (MDS-UPDRS Part III in the OFF state, OFF time, and levodopa-equivalent daily dose) and ^{18}F-DOPA PET, a functional measure of dopamine-synthesis capacity that serves as a biological marker of graft survival, read out at 18 months with a 5-year extension. As a mitigation informed by the first-in-human infection, Phase 2 assigns a dedicated immunologist to each patient.

Competitive context

The field has moved from concept to late-stage clinical and regulatory reality. A Phase 3 trial of a competing dopaminergic cell therapy is underway, and a similar product received conditional approval in Japan in early 2026 [12,13]. These developments validate the modality and raise both the pressure and the payout on execution: manufacturing, scale, usability, cost, and the quality of what reaches the patient will be critical. That is where the platform adds value. We make no comparative efficacy claim for STEM-PD beyond what confirmed, comparable data support.

2.

Own the asset, control the data, close the loop

This section describes a reinforcing, self-improving system. First, the platform can improve STEM-PD by making comparability and scale-up more directed and efficient. Second, the asset can improve the platform by providing critical longitudinal data from production and ultimately clinical validation.

Ownership closes that loop by giving us control over the experiments, assays, timelines, and resulting data. Over time, the same playbook may define what a next-generation STEM-PD product could look like.

Controlling the entire data cycle is the key aspect that ownership uniquely enables and positions us well both in the short term with STEM-PD and beyond.

OUR DATA FLYWHEEL

How we want the platform to improve STEM-PD

By applying our models to STEM-PD, we believe it's possible to reduce the number of experiments required to achieve comparability. This gets us to the clinic more efficiently, saving time and money, and we estimate that reducing experiments by around 30% would save us over two months of development time and over $1M, a goal that is both worthwhile and achievable.

How we want STEM-PD to benefit the platform

We believe that clinical data generated from the STEM-PD program will serve as a critical data stream for both our signaling foundation model and the BioBrain reasoning layer. Longitudinal tracking of bioreactor transcriptomes can directly contribute to our foundation model, and animal and clinical studies measuring engraftment with PET scans can inform BioBrain's reasoning.

2.1 PUTTING THE PLATFORM TO WORK: COMPARABILITY AND SCALE-UP

The most immediate place the platform can provide tangible value is the manufacturing challenge that every cell therapy eventually faces, and that STEM-PD faces now. This is also where the economics live: lowering cost of goods is not only an operational win, but it also may expand the treatable population through better pricing and access and open markets beyond the US, EU5, and Japan.

Today's process grows cells in two dimensions on flat, coated surfaces. It works for early clinical supply and can scale out but does not scale up as we seek to decrease costs and increase market size. Currently, the steps are manual and the vessels are opened during handling, which limits throughput and is hard to square with commercial-scale quality expectations. The path to scale is a closed process, in which cells are grown on microcarriers in a bioreactor that can be enlarged and run as a closed end-to-end system.

THE GOLDEN WINDOW

The timeline for applying our platform to STEM-PD has lined up. Phase 2 is proceeding now on the established 2D product, which protects the clinical timeline. That same period is our window to build and prove the scalable process, so it is ready to carry out the pivotal Phase 3 program [14]. If we do the process work during Phase 2, then Phase 3 can launch on a process that scales; miss it, and the choice narrows to delaying the pivotal trial or running it on a process that cannot reach commercial scale.

Any such process change requires a **comparability** package: evidence that cells from the new process are the same drug as the clinically validated ones, judged across identity, purity, morphology, genetic stability, and potency [15,16,17]. If comparability is proven, the new process inherits the clinical history of the old one. If not, one risks repeating clinical work, meaning years and substantial cost for a cell therapy.

Two of the product's attributes are critical to hold comparable through scale-up and are things that our platform is particularly well suited to inform. **Purity:** cell differentiation protocols can yield mixed populations including off-target cell types, so we want to maintain the current high levels of purity during process improvements. **Potency:** the mechanism of action of the product requires cellular maturation in vivo into a graft. Holding purity and potency steady as the process changes is central to the comparability case. The platform's ability to predict how signals affect cell type transitions makes it a powerful asset in maintaining these attributes.

Part of the evidence in the comparability package is molecular, and so the platform can give insights into which signals may need tuning during scale-up process development to ensure that the cells produced are equivalent. The formal comparability case rests on established release assays, which show that both processes generate the same cells. What the platform adds is the development engine behind that case. We can sample a bioreactor through a port over the course of a run and sequence the cells at single-cell resolution as they develop; the model reads which signaling pathways are active and when and compares that trajectory against the validated 2D reference. The available moves are raw material changes, concentration, timing, and withdrawal — and the model can predict which of those changes would steer the bioreactor cells towards the 2D reference trajectory. A directional answer is already valuable: bioreactor runs are slow and expensive, and narrowing the number of conditions to evaluate is the real bottleneck. **Even a reduction in just three bioreactor experiments would yield estimated cost savings in excess of $1M and reduce experimental time by over two months.** The same reasoning can search for ways to shorten the process or cut reagent use (a direct cost driver) while maintaining comparability.

Our goal is larger than matching a few markers. Ultimately, the scaled-up protocol should recreate the same developmental trajectory, with equivalent maturity and functional potency. Varying process parameters, including conventional ones such as pH, oxygen, temperature, and viscosity, and reading how each impacts cell state, we can fold them into a single predictive framework that connects process, product, and performance, and use it to define a control strategy and design space

Even a reduction in just three bioreactor experiments would yield estimated cost savings in excess of $1M and reduce experimental time by over two months.

A platform improves only as fast as its feedback loop. When we own the program, we decide what gets measured, on what schedule, with which assays, and we keep every result.

in the sense of ICH Q8/Q11 [18,19]. The capsule data engine lets us explore far more of that space than conventional design-of-experiments allows. The goal is a controlled scale-up that reliably meets release specifications and produces a robust and consistent product with lower cost and failure risk. Our near-term path is to evaluate BioBrain predictions first in a small benchtop bioreactor, where directional predictions can be tested cheaply and quickly, and then move to a larger, production-scale system. Note that even once comparability has been established, the platform can still earn its place by optimizing the timing of the process, reducing the amount of reagents it consumes, and reducing the number of GMP run failures. Deliberately perturbing process parameters and reading the consequences also shows how a run can fail, so the critical parameters can be identified, monitored, and controlled. Understanding the full process and robustness levers derisks manufacturing and can themselves become a barrier to entry for potential competitors.

2.2 CLOSING THE LOOP: HOW THE ASSET IMPROVES THE PLATFORM

A machine learning model that is measured by how closely its predictions resemble measured cell states has a blind spot: a cell's profile harbors thousands of variables, but *a priori* it is very hard to know which of those variables affect whether the therapy works. The key is to anchor the loop to a functional outcome: for STEM-PD, the potency metrics [17,20] — and, as the ultimate measure, how well a batch engrafts in nonclinical models [8]. Together, these measures capture what the cells *did*, not how closely they resemble a reference.

Owning STEM-PD is what makes that loop possible, within certain bounds. Notably, we are not building a model to predict clinical outcomes from a handful of early patients, and we are not chasing lot-to-lot variability — every GMP lot is meant to be identical. The loop connects manufacturing, analytics, and nonclinical biology, anchored on potency and graft performance, where controlled experiments can provide definitive answers. Clinical results and imaging remain the program's ultimate validation, and over time can corroborate which product attributes matter most.

These outcomes feed back through a STEM-PD-focused version of BioBrain as a knowledge stream. Knowledge streams are the right way to use these outcomes because the data are sparse and not suitable for conventional machine learning techniques. Clinical cohorts are small and partly observational; they will never have the volume of a perturbation screen. The BioBrain architecture is built to use exactly this kind of input — a small, graded set of real batches with known outcomes, learned as a signature and weighed alongside the literature and the model rather than folded into an underpowered machine learning model of its own. In this context, a handful of well-linked outcomes can be worth far more than it may seem based on numbers alone. These data provide reference points for the model, and they point to the next controlled experiments in product manufacturing.

2.3 WHY OWNERSHIP IS THE RIGHT STRUCTURE

A platform improves only as fast as its feedback loop. When we own the program, we decide what gets measured, on what schedule, with which assays, and we keep every result. Ownership means we own two critical types of

data: the process development that describes how the cells were made, and the clinical data that describe what happened in patients. We are able to ensure that these data are collected in a manner that is maximally informative for our foundation model efforts and don't become siloed, imperfect data streams with limited visibility and utility.

It is important to keep in mind how the relative abundance of data affects how it can be used within our platform. Process development runs, potency assays, and nonclinical models can all be leveraged and combined with the transcriptomic assays that power our foundation model. These data can support the larger numbers required for model training. Clinical data serve as the ultimate validation and can document the spectrum of outcomes, but the numbers are inherently relatively limited. That said, these data can be leveraged as data streams into BioBrain, allowing for our insight layer to reason with these data even if they are not explicitly captured in the signaling foundation model.

An IP moat built from data

Because we generate this data ourselves, much of what we learn is protectable. The role of specific cell types in graft function, potency-prediction methods, and critical process-control steps can be patentable, and the underlying know-how can be held as trade secrets. The data asset is not only a modeling input; it is the foundation of a defensible IP position.

2.4 THE FUTURE: WHAT THE NEXT STEM-PD MIGHT LOOK LIKE

Beyond the target neuron.

A strength the program has under-explored is everything in the graft that is not the target neuron. Dopaminergic neurons are not the only participants in a functional graft: the transplanted neurons wire up much as they would in the intact midbrain, and the non-dopaminergic cells around them — serotonergic and other precursors — can modulate dopamine-neuron function. We may find that a graft carrying dopamine neurons together with the right supporting cells performs better than a purer one; early data from mixing ventral-midbrain and forebrain populations point that way. Our platform's ability to predict cell type transitions and signals that govern their interactions makes it well-placed to help guide these interactions in a precise fashion. Today's lots remain defined by their release specifications, but this is an example of the sorts of questions our synergistic approach would enable.

Our platform's ability to predict cell type transitions and signals that govern their interactions makes it well-placed to help guide these interactions in a precise fashion. Today's lots remain defined by their release specifications, but this is an example of the sorts of questions our synergistic approach would enable.

Toward a next-generation product.

As we learn which cell-state features track potency and graft performance, we sharpen the definition of a functional dopaminergic neuron — and open the chance to specify a better one. The ultimate goal is a developmental trajectory tied to better engraftment and neurite outgrowth, tighter and more reproducible cell-type ratios, faster efficacy, or delivery at lower doses — gains a pure dopaminergic neuron graft cannot reach on its own. Realizing any of them as a second-generation therapy would need a separate program, one integrated from the start with the insights that our platform provides.

3.

Beyond STEM-PD: from one therapy to any cell therapy, and beyond

A platform, practically by definition, generalizes across multiple problems. We see two tiers of generalization, one near and concrete, the other broader and held as optionality.

TIER ONE: ANY CELL THERAPY

The problems we are solving on STEM-PD are not specific to Parkinson's disease. Every cell therapy must scale its manufacturing, demonstrate comparability through process changes, control the identity and potency of a living product, and understand what is actually delivered to the patient. A platform that measurably de-risks those steps is valuable to the entire modality, which is expanding quickly. STEM-PD is the first full-stack demonstration; the manufacturing and comparability capability it proves is reusable across cell therapies, with far less relearning each time.

TIER TWO: ANY CELL STATE

More broadly, the same model of context-dependent signaling supports three kinds of control over a cell's state:

FORM OF CONTROL	WHERE IT APPLIES	CORE QUESTION
Create a state	Differentiation and regenerative medicine	Which sequence of signals moves a cell toward the target identity? [2,3]
Remove a state	Regenerative medicine and cell-therapy manufacturing	Which conditions block the formation of unwanted cell types?
Preserve a state	Cell-therapy manufacturing and delivery	Which conditions hold identity, potency, and viability through scale and administration?

STEM-PD covers two of these: creation and preservation. The proprietary process and outcome data we generate also build a deep understanding of how a dopamine neuron acquires its identity – the question behind creating a state. Also, current work is focused on preserving a state in the sense of refining and holding the identity of a locked product. These same points could be leveraged to specify an improved neuron. We consider this as a forward- and partner-facing optionality rather than a current objective: the near-term experiments serve the existing product, while the understanding they generate could be carried into new programs, whether our own or a partner's.

4.

Validation plan and open work

The following validation gates should anchor scientific and investor diligence.

CLAIM	WHAT WOULD PROVE IT	FAILURE CONDITION
Comparability model	The model prospectively predicts which concentrations, timing, or withdrawals keep the bioreactor process cells as close as possible to the 2D reference, and helps explain product similarity.	No reproducible lift over strong design-of-experiments or mechanistic baselines; material differences remain unexplained.
Process control	Monitoring the critical process parameters the model flags reduces aborted and out-of-spec runs, derisking Phase 3 manufacturing and commercial supply.	Flagged parameters do not predict run failure; no reduction in deviations or scrap.
BioBrain recommendations	A locked recommendation yields a prospective biological result not selected by the standard protocol.	Novel recommendations fail without producing useful discrimination or negative training signal.
Transfer across cell therapies	A pretrained representation improves a second cell-therapy task with less data or better accuracy.	The new task starts from zero and shows no transfer advantage.

The durable asset is not any single model, recommendation, or even STEM-PD itself. It is the owned, closed clinical learning loop with manufacturing and outcomes under one roof, proven first on STEM-PD and built to be reused across cell therapies. That is the system the rest of the field will find hardest to reproduce, and it is the reason to own the asset rather than rent access to one.

5.

References

1. **US Food and Drug Administration.** Potency Tests for Cellular and Gene Therapy Products. Final Guidance for Industry. January 2011. Docket FDA-2008-D-0520. Official source

2. **Rito T, Libby ARG, Demuth M, et al.** Timely TGFβ signalling inhibition induces notochord. Nature. 2025;637:673–682. doi:10.1038/s41586-024-08332-w

3. **Ciruna B, Rossant J.** FGF signaling regulates mesoderm cell fate specification and morphogenetic movement at the primitive streak. Developmental Cell. 2001;1(1):37–49. doi:10.1016/S1534-5807(01)00017-X

4. **Rood JE, Hupalowska A, Regev A.** Toward a foundation model of causal cell and tissue biology with a Perturbation Cell and Tissue Atlas. Cell. 2024;187(17):4520–4545. doi:10.1016/j.cell.2024.07.035

5. **Paul G, Bjartmarz H, Björklund A, et al.** STEM-PD trial protocol: a multi-centre, single-arm, first-in-human, dose-escalation trial, investigating the safety and tolerability of intraputamenal transplantation of human embryonic stem cell-derived dopaminergic cells for Parkinson's disease. BMJ Open. 2025;15(12):e107597. doi:10.1136/bmjopen-2025-107597

6. **Kirkeby A, Nelander J, Hoban DB, et al.** Preclinical quality, safety, and efficacy of a human embryonic stem cell-derived product for treatment of Parkinson's disease, STEM-PD. Cell Stem Cell. 2023;30(10):1299–1314. e9. doi:10.1016/j.stem.2023.08.014

7. **Barker RA, Drouin-Ouellet J, Parmar M.** Cell-based therapies for Parkinson disease—past insights and future potential. Nature Reviews Neurology. 2015;11(9):492–503. doi:10.1038/nrneurol.2015.123

8. **Kirkeby A, Nolbrant S, Tiklová K, et al.** Predictive markers guide differentiation to improve graft outcome in clinical translation of hESC-based therapy for Parkinson's disease. Cell Stem Cell. 2017;20(1):135–148. doi:10.1016/j.stem.2016.09.004

9. **Kriks S, Shim JW, Piao J, et al.** Dopamine neurons derived from human ES cells efficiently engraft in animal models of Parkinson's disease. Nature. 2011;480(7378):547–551. doi:10.1038/nature10648

10. **Paul G, Bjartmarz H, Kirkeby A, et al.** Human embryonic stem cell-derived dopaminergic cells for Parkinson's disease: a phase 1/2 open-label trial. Nature Medicine. Published online July 9, 2026. doi:10.1038/s41591-026-04525-0

11. **Hiller BM, Marmion DJ, Thompson CA, et al.** Optimizing maturity and dose of iPSC-derived dopamine progenitor cell therapy for Parkinson's disease. npj Regenerative Medicine. 2022;7:24. doi:10.1038/s41536-022-00221-y

12. **BlueRock Therapeutics LP.** First Parkinson's disease patient treated in BlueRock's pivotal Phase III trial of investigational cell therapy bemdaneprocel. September 22, 2025. Company release

13. **Sumitomo Pharma Co., Ltd., and RACTHERA Co., Ltd.** Announcement on the Approval for Manufacturing and Marketing Authorization of the Allogeneic iPS Cell-Derived Dopaminergic Neural Progenitor Cell Product AMCHEPRY in Japan. March 6, 2026. Conditional and time-limited approval. Company release

14. **US Food and Drug Administration.** Chemistry, Manufacturing, and Controls Flexibilities for Developing Human Cellular and Gene Therapy Products for a Biologics License Application. Final Guidance for Industry. May 2026. Docket FDA-2026-D-4692. Official source

15. **International Council for Harmonisation.** Q5E: Comparability of Biotechnological/Biological Products Subject to Changes in Their Manufacturing Process. Current Step 4 version. November 18, 2004. Official guideline

16. **US Food and Drug Administration.** Manufacturing Changes and Comparability for Human Cellular and Gene Therapy Products. Draft Guidance for Industry. July 2023. Draft - Not for Implementation. Docket FDA-2023-D-2436. Official source

17. **US Food and Drug Administration.** Potency Assurance for Cellular and Gene Therapy Products. Draft Guidance for Industry. December 2023. Draft - Not for Implementation. Docket FDA-2023-D-4299. Official source

18. **International Council for Harmonisation.** Q8(R2): Pharmaceutical Development. Current Step 4 version. August 2009. Official guideline

19. **International Council for Harmonisation.** Q11: Development and Manufacture of Drug Substances (Chemical Entities and Biotechnological/Biological Entities). Current Step 4 version. May 1, 2012. Official guideline

20. **Bravery CA, Carmen J, Fong T, et al.** Potency assay development for cellular therapy products: an ISCT review of the requirements and experiences in the industry. Cytotherapy. 2013;15(1):9–19. doi:10.1016/j.jcyt.2012.10.008