Can Breakthrough Innovations Be Made Systematically? A Conversation with Noubar Afeyan

By Gary Pisano

Noubar Afeyan has dedicated his career to improving the human condition by systematically creating science-based innovations that are the foundations for startup companies. At Flagship Pioneering, which he founded in 2000, Afeyan built an enterprise where entrepreneurially minded scientists ask, “What if?” and iterate toward the answer “It turns out …” in order to create first-in-category companies in health and sustainablity. Over two decades, Flagship has fostered the development of more than 100 scientific ventures resulting in $30 billion in aggregate value, thousands of patents and patent applications, and more than 50 drugs in clinical development.

Afeyan spoke to Gary Pisano, a professor at Harvard Business School and the author of *Creative Construction: The DNA of Sustained Innovation,* at Flagship's offices in December 2019.

Gary Pisano: What was your journey from the time you were young to starting Flagship?

Noubar Afeyan: I was born to an Armenian family living in Beirut, Lebanon, and in 1975 immigrated in a hurry to Canada, which graciously accepted my family as escapees from the civil war, giving us citizenship. I grew up in Canada, went to college there, and ended up coming to Boston to attend MIT. My interests at the time were in engineering, particularly chemical engineering. But the frontier of that field was focusing, for the first time, on biological engineering. So I did my PhD in one of the very first cohorts to be educated in this hybrid way between engineering and science.

In 1987, I ended up taking a big leap of faith and starting a company, which back then was not the traditional path. I grew and ran PerSeptive Biosystems for 10 years, five of them as a public company. It ended up becoming the largest instrumentation company on the protein side of the life sciences. I got involved in a number of other startups along the way in the '90s. That all led to the formation of Flagship some 20 years ago.

GP: Why Flagship? Why didn't you just start other companies? Some folks start a company and then they start another, but you took a different route.

NA: PerSeptive Biosystems was an intense 10-year, postgraduate education in everything related to applying science, engineering, innovation, IP, and marketing to business. It was an intense learning process. One of the things I became drawn to was all the ways in which you could create value from breakthroughs. Because I had the good fortune of having a strong team of innovators inside a company of some 880 people, over 200 with advanced degrees, I became interested in doing entrepreneurship in parallel. I was running one company, and I started spinning out other companies and partnering to start new companies. I realized that becoming a co-founder is an interesting concept relative to a solo founder and that you could do more than one thing in parallel. The notion of parallel entrepreneurship made me ask: Can you do more this way than improvisational entrepreneurship? That led to the founding of Flagship.

In 1997, PerSeptive successfully merged with the DNA leader in the instrumentation space [Perkin Elmer/Applera, in a stock swap valued at $360 million]. A number of the other startups I was involved in co-founding had gone public, or had been sold. And on the heels of all that, I developed the basis of the belief that the only way to do this type of activity for a living was not serial entrepreneurship, which is the traditional notion, but a more disciplined professional process. I set out on a journey: Flagship is just the instantiation of figuring out whether, how, how well, with whom, one can do this.

In its initial formative stage, Flagship was called NewCoGen, which stood for “new company generation.” Two years later, people had convinced me that it sounded like a disease of some sort, and I picked a different name. But new-co-gen was what we did. We wanted our products to be other companies, and that journey continues 20 years later with exactly the same notion, although we’re getting a little better at it and have scaled quite a bit.

It’s a completely different way to think about entrepreneurship, because the context within which startups operate historically, at least when I started, celebrated a romantic, chaotic, improvisational kind of culture. They tell majestic birthing stories, where eventually, if the company succeeds, the birthing process gets more and more glamorized. And the reality is it’s messy, unpredictable, fraught with risks and errors. And whether it can only be done that way instead of being a little more thoughtful, a little more planned and disciplined, and a little more consequential, seemed like appropriate questions to ask, especially having received a classical engineering education. It’s been an interesting battle to figure out what can and what cannot be operated in an institutional way when it comes to startups.

GP: What's different about Flagship from a traditional venture capital firm? You do raise funds. But from what you're describing, you're not a venture capital firm at all.

NA: We're not actually structured as a venture capital firm. Or put it this way: A venture capital firm is just a part of how we're structured. It's how we source capital: We raise funds and deploy them to finance our activities over a limited time frame. Instead, we are organized as a company where our deeply experienced leadership team manages all aspects of company creation and development in the same way that another business would manage the process of making medicines, or cars, or what have you.

So we centrally have operations that involve research and development, which we call origination; growth, which is like manufacturing; and finance, IP, legal, etc. But we operate organizationally as a company that happens to avail itself of venture financing. We don't have a singlular consolidated balance sheet as would a conglomerate. We're not operating off a permanent base of capital. We're going through cycles of fundraising with each representing a new beginning. Part of the reason is historical: When we started, the only money available for this kind of portfolio came from the people who invested in other portfolios of companies, who invested in venture capital.

But more specifically, the way we're different is the following: In traditional venture capital, everything that happens involves science that comes from academia or research institutions, hospitals, etc. If you bring in scientific breakthroughs that have already been made, hire appropriate people, and deploy capital, you can do a startup today in the life sciences. By the way, in each of those categories, there are multiple instances of each. So there are multiple sources of science in typical venture-capital funded startups, multiple people that get brought together, and multiple sources of capital. It's all syndicated.

"I realized that becoming a co-founder is an interesting concept relative to a solo founder and that you could do more than one thing in parallel. That led to the founding of Flagship."

All those things must come together and coexist and operate harmoniously through thick and thin and create value: That is a successful startup. In my 32 years of starting companies, and a lot of it with venture capital, I've come to the conclusion that the most remarkable thing is that this method actually works once in a while. The science may not work, or the products may not work; the organization is fraught with all sorts of conflicts and all sorts of misinformation. What Flagship set out to do, and what has evolved over the last decade at a much larger scale, is to combine the capacity to innovate, the people to lead the formation of companies, the people to grow companies, and the capital, all in a singular institution.

So Flagship is completely different from a capital provider in that we're not providing capital to anybody. We are deploying our own capital around our own ideas and our own team to generate companies that we eventually spin out. There's no notion of spin-out in the venture world. There's no notion of having all the elements win together or lose together, as opposed to having completely different agendas. All of those things are absent. Of course, we have our own additional set of risks. But first and foremost, organizationally, Flagship is very much an innovation company, which accesses capital.

GP: You manufacture firms. That's the way I've come to understand it—or you manufacture new ventures.

NA: We produce new ventures. And in fact, the shares that we ultimately sell to others who co-invest with us are their options to the future value created by those companies. That's the way we view it.

GP: You've used the phrase "institutionalized entrepreneurship" to describe what you do. That's interesting, because for some people that combination of words is an oxymoron.

NA: Since I have the opportunity, let me repeat what I usually say about entrepreneurship. First, it's ironic that the English word's root is a French word for undertaking, so the literal translation of "entrepreneur" is an undertaker—which is funny because usually entrepreneurs preside over things that don't exactly go well. Secondly, and more seriously, I think the word is a misnomer in that the suffix "-ship" in English means the state of being something, its condition. We talk about leadership, sportsmanship, friendship, or craftsmanship: These are all states of being something. Well, you know, I'm an engineer. I don't do engineership; I do engineering. That ending—entrepreneur*ing*, if you will—would be much more appropriate to how I think about the activities involved in creating companies, rather than just the state of being an entrepreneur.

By the way, this emphasis on *being something* is not much different from architecture or medicine before they became real professions. People used to think that people were born with a predisposition to be an architect or doctor, as opposed to learning, perfecting, improving the practice of architecture or medicine. When people say of entrepreneuring, "Yeah, but you can't create a Bill Gates," the answer is, that's true. You also can't create a phenomenal brain surgeon. That doesn't mean that people can't do brain surgery. People can learn how to do brain surgery, and the phenomenal ones will always be phenomenal. It's a complete misunderstanding to think that the only way you can be a phenomenal brain surgeon is to have no preparation and wing it.

This sort of thinking bothered me in my 30s. And I decided to do something about it in my 40s, and now in my 50s I'm increasingly certain that entrepreneurship can be done institutionally. I wish the world would call it "entrepreneuring," but so far I can't seem to change the language.

You can do entrepreneuring individually. A founder is just the conductor of a process: Founders pull together the orchestra; maybe they play an instrument or two themselves; but they cannot do it all on their own. The person who organizes the resources, sets the direction, and undertakes the most risk—more than anything else—is the founder. But you could also think of those functions as the purpose of an organized team, occurring within a learning environment where entrepreneuring is done collectively and practiced repeatedly against metrics that increase what leverage can be brought to bear. Done individually, collective action, serial company creation, and measurement and leverage are a distraction at best. So by saying "institutional entrepreneuring," I'm really comparing it to institutional investing.

"What Flagship set out to do is to combine the capacity to innovate, the people to lead the formation of companies, the people to grow companies, and the capital, all in a singular institution."

When I was an entrepreneur for the first time, starting PerSeptive at 24 years old, I looked at venture capital firms and I was very envious. All of them seemed to know an incredible amount. They were all geniuses to me. They could look at their Rolodexes and contact 50 different CEOs and 10 different partner companies. And I thought, how could they know all this? By comparison, we entrepreneurs knew nothing. And the answer was, because venture capital firms are institutions. They have institutional memory. They have a set of people who invest in companies over and over again. People leave, and other people come. They think about the institutionalization of investment over the long term, particularly in the world of early-stage companies. I thought, why can't one also do that with how innovations are made and how the companies are formed?

And of course I was told, "No, no, no, no." And the reason was that investors wanted to build portfolios and the individual constituents of their portfolios had to be entrepreneurs that had no other activity. So the notion was, you can only diversify if the pieces in it are not diversified, so that every one of them was in this life-or-death struggle. And I thought, well, how is that fair for the entrepreneurs, the teams, and the science, that it has to be completely boxed in to one out of 20 bets per fund cycle? That was my initial thought, 20 years ago.

Gradually, what Flagship has begun to do by institutionalizing this act of innovating and entrepreneuring is to say, can't all of it be part of portfolio formation? So that as an entrepreneur, if you're more inventive or better at strategy, you can contribute to many company foundations and still be part of dedicated teams working as part of a single institution. That's the experiment we're running.

GP: The traditional model of venture capital as you've described it creates all sorts of conflicts of interest. It creates what we economists call moral hazard. You've got an entrepreneur who is completely undiversified in their career risk, whose financial risk is totally undiversified. But the venture capitalist has a diversified portfolio: They're hedging their bets. That leads to some dysfunctional behavior, doesn't it?

NA: It certainly creates a lack of trust and alignment. It's hard to align when your end result is going to be different. Ironically, beyond the risk you run as an entrepreneur, you want the company to work and you're supposed to be all in, and therefore you will do whatever you can with the information you've got to survive. And it's not clear that is necessarily the best position to put an entrepreneur or a team in if you want to make sure that they're asking the right questions in the right sequence.

But an investor has no choice, because investors want entrepreneurs in that box. The reason, by the way, serial entrepreneurs tend to do better, from the data I've seen, is that they learn the value of diversification, and their learning cycles help avoid getting completely pigeonholed into some arbitrary idea that they now have to take forward. This struggle between startups, science, the sources of science in academia, and venture capital is very clear to most people. But they think it's all part of what it takes to succeed. Flagship is asking, is this the only way to do it?

GP: Describe Flagship's process of company creation in detail.

NA: Maybe I should start by saying that a big part of what we do is highly iterative, fashioned after the idea of Darwinian evolution. We experiment with how we do what we do, all the time. It's an integral part of the culture to question and to adjust—sometimes improve, sometimes not.

Over time, we decided that applying what we do, with the rigor with which we apply it, to a crowded space, to opportunities that everybody can see because they are adjacent to what is already being done, would not be the best place to situate ourselves—even if we could more professionally conceive and create companies that have certain advantages over others. We decided a decade ago to move away from what is here and now in order to work on what others might view as distant, future-oriented projects. Our process is designed to work in that area code: on the kinds of things that people think five years from now might sound really reasonable but today have no connection to reality. We also reasoned that the notion that we would wing it, in such a way-out-there place, was reckless. So we said, either work on well-known spaces in an improvisational way; or, if we're going to go out to completely unknown places, we should at least protect ourselves with a repeatable process. So we thought a lot about what could be repeatable. For where we aim to operate, it's a form of protection.

We have a four-phase process. It starts out with what we call explorations. Explorations are conducted by small teams that generate what we call venture or value hypotheses. Explorations start out by wondering about a new space that *might* exist. We don't have any reason to believe it does. We ask questions like, "What if this could be done?" or "If only that could be done, what would that mean?" By asking those kinds of questions directionally, an exploration takes that initiative and instantiates it into a set of hypotheses. What is a venture hypothesis? It is a scientific or technological advance that has not yet been made but that we assume could be made. That results in an imagined product, sometimes a service, that does not yet exist, which can deliver value to a putative beneficiary that of course we're also describing but who doesn't yet know they could have this capability, let alone the need. And that allows us to describe a hypothetical future source of value. We don't do one of these; we do several. Why? Because we want to make sure we don't have the bias of our starting point. Once you start at one place, our fear is that you'll end up in proximity to that initial place. We want to sample a broad set of starting points within a particular zone that we're after.

Next, we take these venture hypotheses and engage the current experts, nonexperts in academia, in startups, in the incumbent large companies, and we present to them our initially bad ideas. I say "bad ideas" because they're uniformly bad—because how could they be good since they're just a starting point? But then by engaging with people who know what there is to be known about a subject and can speculate a bit about the future, we get massive critical feedback. And in fact, that's part of the scientific process of organized skepticism. We tap into a collective source of organized skepticism through a broad network of people. And what happens is they engage with these ideas and they tell us all the things that were wrong with them. Some of them are knockout blows. They tell us, here's exactly why it can't work. Maybe it violates some law of physics. Perhaps it's already been done. There's a set of things that will get thrown at us and our team. Then our teams iterate the hypotheses to overcome those problems if possible, or else say, okay, that's a dead end—let's move on to another one.

So what happens is that descendant versions of the original conception are constantly being created as we're trying to overcome the objections that are coming to us, as best we can. We do some 50 to 100 of these types of Explorations per year, so that there's a scale to this process. We're not hoping that the one thing we're working on produces a company. We're actually quite ambivalent as to which company comes out of which starting point or which exploration. It's very much a merit-based thing. But if a hypothesis, or a descendant of a descendant of a hypothesis that we started with in the beginning, becomes, as we present it to more and more people, something they cannot find fault with, then we go to the next phase.

GP: There's no lab work at this stage?

NA: There's no lab work yet. In fact, our firm belief is that there's no point in showing that something can be done if you can't first understand for whom it's useful and to what extent. So we purposely do not look at current science or find applications for it. That's what the rest of the world does; that's what innovating within adjacencies is all about. And they're good at it. By the way, some adjacencies can be far out. Others could be closer in. But the reason they're adjacencies is that they're grounded in the present. We're trying not to operate in a way that's grounded in the present. We're trying to say, "Okay, what *could* exist?" And then we want to work backward. That's what the second phase is about.

"We're trying not to operate in a way that's grounded in the present. We're trying to say, 'Okay, what could exist?' And then we work backward."

What we call the prototyping phase is very much copied from the way industrial companies design a product: First you do research and development; then you put something together against some list of goals for the product. But then you must actually prototype it, and you see whether the thing works and you beat it up, and you show it to people and they beat it up. So experimentally prototyping the venture hypotheses is our second phase, and we call the result ProtoCos. We made up the word to describe not a company but a prototype of a company. It's something that wants to become a company but isn't yet a company. In that phase, it's the founding team who worked on the exploration, plus maybe three or four people that have expertise in a particular scientific area. This team of maybe six or seven will engage in answering literally killer questions. We want to falsify our hypotheses experimentally.

Interestingly, this notion of experimentally verifying our hypotheses is the same idea as reduction to practice for an invention. In other words, underlying these things is an invention that is looking to be made by the experimental validation of a hypothesis. If we get that validation in the experimental phase, it typically takes nine months to 12 months of work. It usually costs us $1 million to $1.5 million to do that work—

which, by the way, very strictly confines the kind of things we can work on because they have to be things that are experimentally verifiable. With that as a limitation, we're not going to work on laser fusion. We're not going to work on space flight, because it'll take $100 million to verify anything.

We're trying to figure out: Can we get some key killer experiments done? The ProtoCo phase will kill off a subset of the seemingly infallible descendant hypotheses that now are being really tested experimentally. Now, one might say, listening to this, "Well, how do you know that the questions that were asked were the best questions and that you won't find out later that it's not going to work?" We don't. But, boy, do we front-load the killer questions, because the culture here is such that people are incentivized to continually wonder whether this is the best thing for them to be working on. And we want them to be critical in that way. The best way to find out if it's the best thing or not is to try to knock it out. So everybody's incentives are to figure out what's wrong with it, versus figure out how to convince the world not to ask the tough questions so that it can survive—because we're kidding ourselves if we do that. There's nobody else to kid. In the first two, three years, it's just us working on the projects.

GP: What's next?

NA: The third phase, what we call the NewCo, is what the world would consider a startup company. By this time, we've spent a year or more on the idea. We start getting the larger team assembled. The exploration team continues to work on the project. One of our origination partners will lead the effort. We start bringing on board scientific leadership. We have a team internally that supports the operations of the company. So all aspects from talent to finance to everything else is done and owned centrally. It's specifically not a service providing model. We do this so that the founding pioneering team can focus on science and products. NewCo phases usually last two years, plus or minus. Some 25-plus million dollars are spent on each NewCo, and it's all financed internally through our own sources of capital.

At which point, if all goes well, the company is spun out with an externally sourced CEO, a board, a leadership team that now has quite a bit of velocity, quite a bit of early de-risking, and a clearer articulation of the products that derive from the original innovation. At this final stage, called GrowthCo, the process is focused on both execution and further pioneering, doing *new* innovations. Our GrowthCos represent a very different stage of a company than traditional startups when you have absolutely nothing and you're trying to put the pieces together. These then are the four phases: Exploration, ProtoCo, NewCo, and GrowthCo.

GP: Can you give us an example of a company that went through this process?

NA: We do maybe six or seven of these a year. And so there are many that exist. They follow the same pattern. So I can pick from probably 60 of them that we could use. They have more in common than they have differences, because we really are operating against this model consistently. Our decision-making process going from an exploration to a ProtoCo and then a NewCo is a stage-gated process. This is ironic because everybody who works in large companies gets enslaved by stage-gate processes and thinks that when they go to startups they're free to not follow any process. Here, that's not the case, because we realize that winging-it is a difficult thing to scale up.

A good example that is now in the GrowthCo phase is the company called Rubius [NASDAQ: RUBY]. Rubius started out as an exploration in 2013, where we wondered whether we could render red blood cells into allogenic, off-the-shelf therapies, because red blood cells are the carriers of oxygen. Of course, we were aware at the time that people were playing with other kinds of cells—for instance, T cells and other immune cells. For a variety of reasons, we did an exploration, asking how we could think of red cells as a therapeutic. We looked pretty systematically. Could we ultrasonically shock things into red cells so they go inside the cells? Could we link things to their surface? Could we convert them into a therapeutic all inside the body? We looked at a whole lot of different ways of doing these things.

Among the ways that we hypothesized might work was not to work with red cells themselves but to work on their precursors. These are what's called hematopoietic stem cells. We were aware of academic work funded by DARPA to produce red blood cells from hematopoietic stem cells. This seemed to be an artificial way to make more red blood cells. Of course, the quantity of red blood cells we would need in their more canonical use as oxygen carriers is gigantic, and that's why it's really never been a cost-effective approach. But we weren't interested in that. What we were interested in, what we hypothesized, was: If you could start with a hematopoietic stem cell outside the body and convert that cell into a red blood cell—which it had just begun to be shown was possible; no rigorous process, no scale, nothing—well, what if we came in with a virus, early in the maturation of that cell into the red blood cell and had the virus now express a protein of interest, or more than one, that in the fullness of the making of the red blood cell would effectively be eliminated, leaving behind only the red cell and the protein or proteins that you want?

That was Rubius's initial hypothesis. There was no scientific basis to believe that it could work. We talked to lots of experts who told us all the reasons why it wouldn't. In fact, there wasn't much known about one of the key processes in making a red blood cell, which is called enucleation. There was also a big mystery about whether you could virally tranduce these cells. The answer to that turned out to be yes. With that as the resultant descendent hypothesis, we went into labs and said, "Okay, if you can do it, you should be able to go in with many different payloads and watch what happens experimentally." We tried some 50 payloads in this way.

Also, you want to convince yourself that you can make enough of the actual protein in any given cell for it to have a therapeutic effect. We reasoned that if you could have, say, 100 copies of a protein per cell and you need a million copies, then you need to give your patient 10,000 cells to get the same job done. I'm just making up numbers, but if you had 100,000 copies per cell, then one cell would do the job, let's just say. By that logic, we wanted to see how much expression we could really get. Because that would dictate whether we need to give patients pints of blood or a milliliter of blood. And of course, a milliliter would be better for both safety and convenience, but also because of manufacturing, because we were going to have to manufacture that much more to get the same therapeutic effect. We had to convince a lot of people that a red cell could be a therapeutic. But we likened it to T cells, which by 2014 already had 15 startups. Long story short, we went from an exploration to a ProtoCo.

During the ProtoCo stage, we technically demonstrated both that you could actually make many different payloads and cells, at least in the first instance, and that you could make quite decent quantities of them. We had reason to believe we could improve it even thereafter. That became a NewCo initially under the leadership of Avak Kahvejian, who is one of our origination partners, and that grew to be some 17 people in the labs. Around this time, we attracted the former CEO of Novartis Pharmaceuticals, David Epstein, who joined Flagship as an executive partner to take on the executive chairman role at Rubius. And Torben Straight Nissen, who came out of Pfizer, joined as president. The two of them led the GrowthCo spin-out phase of the company, attracted quite a bit of capital, grew the team, and just a year and a half ago, we added as CEO Pablo Cagnoni, who was from Novartis and Amgen and had a long, distinguished background in both drug development and scientific medicine. The company is now just about to dose patients and has multiple programs developed only a few years after our initial notion that was based on no real experimental proof of concept.

GP: I wrote a Harvard Business School case study about Flagship. When I teach the case, it always strikes me that the process *sounds* so easy. And yet there's something here in terms of the behaviors that is absolutely critical to make this institutionalized process work. Talk about Flagship's culture.

NA: There are several elements that come to mind. First of all, you know, we are living at a time when entrepreneurs are celebrities. People who work here are not celebrities, and they understand that what they're doing is a responsible act. We're taking other people's money. We are causing people to leave their jobs and join projects we've proposed. And there is a seriousness to that, you know.

"People talk a lot about celebrating failure, as if it were an ancient saying by a sage. But in our case, failure is integral to what we do. It's celebrated here because we know that if we fail more, we will find some valuable things."

Another element of our culture is that we accept that breakthroughs and things that are transformative *emerge*. They're not designed or specified a priori. We have reconciled ourselves to the fact that we are managing an emergent process. I spoke earlier about Darwinian evolution. Darwinian evolution, simplistically, is variation, selection, and iterative cycles of descendancy or inheritance. It turns out we know about that in DNA and in nature. But you can look at the same phenomenon in ideas and products and all sorts of spaces and you see the same variation, selection, and iteration. It's not what we're brought up with in our schools. We don't actually expose ourselves to problem-solving in this way. When we hire our associates for Flagship's origination teams, we hire out of a graduate education program. We haven't so far tried to rewire or retrain people to figure this out after they've had a lot of experience, because what would have made them successful in the more traditional goal-based way of doing innovation would be lost at Flagship. So the culture at Flagship is one where you must subject your ideas to fierce criticism and not have that depress you.

Relatedly, to be able to articulate initial ideas that you know are kind of silly in order to arrive at serious-seeming ideas, is super difficult to behaviorally adjust to when you come out of a school where you're the top in an area and you've been rewarded for being right. And when you get here you have to say things that are almost certainly wrong or not known to be right all the time in order to discover the ones that are actually valuable. In the fullness of time, I think that transformation of people's thinking can be achieved.

Those are all things that are based on rewards, based on just how we think about what matters. People talk a lot about celebrating failure, as if it were an ancient saying by a sage. But in our case, failure is, in an evolutionary sense, the discriminator that produces success. The reality is, failure is integral to what we do. It's celebrated here because we know that if we fail more, we will eventually find some valuable things.

GP: What do you do as a leader to help shape this culture?

NA: One of the mindsets that I feel is important in general in entrepreneurship, but essential to the form of extreme entrepreneurship we practice at Flagship, is what I'd call a kind of paranoid optimism. If you're only optimistic, which most people who do entrepreneuring are, you will do things that have no checks and balances. If you're only paranoid, you're generally depressed because you don't think anything is going to work. Being willing to toggle between them is an important capability. Well, I embody that because I'm never taking what we're saying for granted or to be the truth. I have a belief that if we can figure out how to do this institutionally and we can let the values that matter emerge out of doing them, we can find people who can embody those values.

"The reason we use the word 'pioneering' is that there's a first-in-kind character to what we do. There are always new things to explore and new places to go. Opportunities will constantly be created for what can come next."

GP: Do you believe machine learning will change how biology is done?

NA: Machine learning is broadly based on what people call artificial intelligence—or, rather, a subset of the field. Ironically, for me, some 37 years ago, the one and only college newspaper article I ever wrote at the *McGill Daily* was on AI and artificial intelligence, circa 1982. And that was fascinating then, because people were starting to use

computers to create programs that were using heuristics the way humans do when they think about certain kinds of if-then relationships: If I see this, I'm doing that. Computer scientists tried to encode that to create what used to be called expert systems. Those were just the beginnings of mimicking human intelligence. Today, of course, much more data is available and is processed on much more powerful computers. Machine learning is extending that in new ways.
What's possible? Three years ago, we started explorations that were centered around capabilities that either were just becoming available or we thought would soon become available, so that we could anticipate them and say, What if? Some six different projects have been born out of explorations in Flagship Labs that are now at the ProtoCo, NewCo, or GrowthCo phase. There's a company called Integral Health that's accessing massive amounts of data to drive clinical decision-making and clinical trials, not as a service but as a platform around which we're operating.

We just announced recently the launch of a new company, Cellarity, which is using machine learning to characterize very specific cell behaviors in disease or as a result of certain drugs. One way to think about these behaviors is as alternative states that cells can be in at a given time and that in many cases may not have been previously described. You can't characterize 20,000 data points per cell in a million cells in any other way, and statistics doesn't really help you at the level of two or three variables, because there are 20,000 variables here. But it turns out that you can train algorithms to detect that these cells are exhibiting similar behaviors, or shifts from one behavior to another. Yet another of our companies is developing some interesting ways to purely computationally generate totally new, unprecedented proteins and DNA sequences with little or no experimental work.
So all of that has gotten us more seriously thinking about what this could mean. When we look at games, people are marveling at the fact that deep learning algorithms can play games that they've never played before. The machines have never played a human before, but they can beat them with algorithms that humans don't understand and probably can't understand. What's really intriguing to us, more than just the

application of this machine learning, is to ask what novelty such a paradigm can generate. If you think about an algorithm that learns and that can generate something completely new, that humans don't understand in terms of how it's doing what it's doing—whether it's a game, or a new product, or a new way to land airplanes, or a new way of diagnosing a disease—guess what? The resulting black box looks a lot to us like biology, because we also don't understand how biology does what it does with all these complex parts that nevertheless produce a result. And so I think there's a special place for the kind of innovations that we've done to apply to the result of machine-learning-generated novelty. That's what interests us deeply for the future, well into the future.

There are other applications of machine learning that we're attracted to: alternative ways to think about health care, because rather than spending a lot of money on therapeutics and surgery when somebody gets really sick, society might decide to spend less money much earlier when people are not sick. To the extent that dollars can be spent early, whether it's through vaccines, early detection, protection, prevention, or delay of disease, we think that will benefit both patients and society—and, ultimately, be where much of health care moves.

GP: If you take the view that Flagship is just a company and its products are other companies, how big can a company like that get?

NA: I have not thought about whether there are natural limitations to this kind of enterprise. I would also say that making Flagship big is not among our top priorities. Sure, we recognize that if scale achieves objectives, we should not be afraid of it. But achieving scale for the sake of scale isn't something that we aspire to do. Now, in the natural course of the development of Flagship, at some point either we or others might well try to do this in another geographic or scientific environment. Of course, that interests me.

I would say just one last thing. The reason we use the word "pioneering" is that there's a first-in-kind character to what we do. That potentially makes it scale-free to the extent that there are always new things to explore and new places to go. Opportunities will constantly be created for what can come next. But the resources that society dedicates to unprecedented innovation will probably always be a small subset of the total. To say it in a different way, I think that people will always think that adjacencies are a safer bet.

Editor's note: *This conversation has been edited for length and clarity.*