NATIONAL ACADEMIES Sciences Engineering Medicine

In-depth Discussion with Anastasia Gamick of Convergent Research

Anastasia Gamick is the Co-Founder and Chief Operating Officer of Convergent Research. There, she pioneers the management of a new model for large-scale science projects called Focused Research Organizations (FROs), which address neglected bottlenecks in biomedicine, biosecurity, climate technology, and other areas. Anastasia was the first operations hire at Neuralink, served as Chief of Staff at the robotics company Creator, and oversaw the Give Directly relationship at the fintech company Segovia. Notably, she also led the scaling-up of COVID-19 test production at Curative, Inc in the first weeks of the pandemic.

Gamick participated in the panel: *Investment Strategies That Will Rapidly Advance Innovation to Markets* during GUIRR's June meeting. The panel discussed how funding models have evolved from Federal government dominance to the growing importance of traditional philanthropy, value investing, venture philanthropy, and new economic models for underwriting research at universities. Here Anastasia shares with the GUIRR community further insight into her views and work to advance the U.S. research enterprise through the novel Focused Research Organizations model.

Your career spans various impactful roles in organizations like Neuralink, Creator, Segovia, and Curative, Inc. How have these experiences shaped your approach to creating and managing Focused Research Organizations (FROs) at Convergent Research?

I've been an early team member in several highly technical, fast-moving startups. I've had the good fortune of watching brilliant and dedicated entrepreneurs solve extremely difficult technical challenges with skill and aplomb. In each situation, there was a clear goal like "build a software tool that could deliver aid payments to hundreds of thousands of beneficiaries in the most remote of Sub Saharan Africa each month" or "scale up the COVID-19 testing capacity of California" or "create an implantable brain computer interface with X input/output channels."

In each of these cases, the leader had to solicit feedback from the team and markets to get the approach right, and they had to coordinate effort across multiple teams from different disciplines – hardware, software, medical, supply chain, regulatory, sales, fundraising and more. The amount that was accomplished in extremely short periods of time in these industrial technical and operational moonshots was magical to watch. I feel fortunate to be able to bring those experiences to more basic science.

At Convergent Research, we launched six FROs in just three years, and so the ability to work at a rapid pace has transferred over. We and our FROs needed to quickly recruit and vet new team members. This is a very different style of recruiting than academic labs tend to do. We needed to create budgets, financial forecasting and operational workflows rapidly, including templates for the FROs to follow. But most of all we need to work with diverse personalities and experts to help them deal with challenges and align and gel their teams – it's all about accessing expertise, and setting up well defined roles, touch points and communication pathways both inside and outside the relevant companies. Because FROs are a new model, and very different from academic work styles in some ways, my ability to rapidly prototype such systems and work under pressure has been invaluable, as has my experience in translating such systems into frameworks that work for highly technical people.

More broadly, over my career, I've evolved towards a more deliberate approach to scale-sensitive impact. I started out with automating food production, progressed to fintech and pandemic response, and as a co-founder of Convergent Research, I'm now applying that counterfactual impact thinking towards the creation of platforms for science. That is reflected both in our choice of ideas and in how we think about what is practical to execute.

Incentivization is a topic that you and fellow panelists addressed in depth. How do you think incentives need to change to encourage more researchers to tackle high-risk neglected research problems?

For scientists who want to participate in a traditional academic pathway, they must demonstrate an ability to publish first-author papers, shape their careers in a way that attracts individual grant funding, and they often need to make sure that their projects are chosen to set up their students for the same academic success.

Sometimes, these pathways are aligned with the most transformational work needed in a field -- and sometimes they are not. If funding and job security is rewarded for a very particular shape of success, the smart players will choose to work on the problems that they know will generate good results in that framework. Again, despite some structural problems in the system, the academic framework works well for many research problems but not all.

Risk is a complex concept. Something might be high risk in academia or industry but low risk with the right team incentivized in the right way. We don't think of risk as one single variable but separate out team and leadership risk, execution risk, technical risk, fundraising risk, post-FRO transition traction risk, and other types of risk.

Let's start with technical risk. Even when building tools for basic science, FROs tend to focus on more engineering and system-building-oriented problems where there is a roadmap and a set of contingencies, rather than on totally unpredictable discovery. Still, necessarily, tackling research problems at a fundamental level means that there is sometimes a meaningful probability of outright technical failures. For that case, we'd try to have some offramps – and increasingly we aim to evolve our model so that we can support more "seedling" projects that de-risk the most basic assumptions behind a FRO with a smaller initial team.

For the risk of taking on such an ambitious, large-scale, cross-disciplinary project, or the risk of going for a project that prioritizes field wide impact rather than individual publications or credit – that's a risk we think isn't fundamental and is one that we can make a big dent in. We are creating an alternative path. This isn't something we can do alone. It requires more government agencies and philanthropists to get involved in building this ecosystem, and it does require pioneering researchers to be guinea pigs for early FROs, as well as for other structures emerging like new private or public ARPA style agencies which need technical leaders who can drive coordinated research programs.

What we've found is that even publicizing that FROs are a possible path, and engaging in conversations with researchers about this model, leads them to think differently. Sometimes they wouldn't even spend the time to think about possible FROs if there wasn't a way to do and fund them. Changing that thinking is part of what breaks the lock in of a single incentive system and lets the community start to move to a more diverse set of incentive systems with which to solve problems.

What strategies have you found most effective in motivating researchers and organizations to work together on large-scale science projects?

Scientists are already motivated, but the systems don't always incentivize it. It's more about removing barriers that prevent scientists from taking moves that might seem like career or professional risks. When there are pathways and possible rewards, many scientists naturally dream big.

Free-wheeling and fluid academic collaboration that fades in and out with individual interest is a lifeblood of science and great strengths of academia. But it doesn't always work well for systems engineering. Nobel Prize winner Eric Betzig said something in a recent piece about a specific problem Betzig wants to solve in compute intensive and hardware engineering intensive live tissue imaging. He said* "This is difficult to achieve in academia, where single-minded focus is rarely possible, resources are often constrained, staff turnover is high, participant incentives are mismatched, security breeds complacency, collaboration dilutes accountability, and consensus fosters mediocrity." That's an extreme case, but in general there are different strengths for different approaches relative to different problems.

The ARPA model, where a full-time program manager coordinates distributed collaboration can

be powerful. This model sets clear milestones, a social scaffold and sense of social pressure, and helping create the necessary exchange and arbitrage across different groups. Without that, some academic collaborations do lead to a diffusion of responsibility or tend to fizzle or become somewhat random assortments of work that each group would otherwise have done anyway. The FRO model works best where extremely high coordination is needed, such that much of the effort needs to be done under full time operational leadership by a single core team, plus some subcontractors and collaborators of course.

Many scientists in more fundamental research fields are still not highly aware of either the ARPA model or the FRO model. By demonstrating success, we can change the thinking about the levels of coordination that are necessary and possible for researchers to engage in.

Looking ahead, what do you see as the biggest opportunities and challenges for Convergent Research to foster a more innovative and collaborative research ecosystem?

I'm excited that so many of Convergent Research's FROs are created with input from the scientific communities that they represent. As each FRO can create a tool or dataset that solves a bottleneck in a particular field, they'll be open sourced, published, or released back into the research environment. This means that FROs will collaborate heavily with a wide variety of academic and research institutes.

We'd like to achieve a scale of impact well beyond the FRO. This could take the form of impact exits like the FRO technology being used in large-scale big science programs, leading to major spinoff companies, or becoming ubiquitously used by a field as open-source products.

This is also a challenge because it means that FROs face a set of entrepreneurial and partnerships tasks as well as pure science tasks, to gain traction and support for what they build post-FRO. This is inherently a multi-institutional, multi-sector, collaborative, and partnership-based proposition. As a challenge in this, FROs are still a new model. We hope to continue to partner with philanthropies, government funding agencies and corporate labs -- but as a new model, not everyone will have funding programs set up to support such transformational work. This means that there will be intense work with all partners to educate and figure out how to best work on funding transformative science together.

At a more macro level, we think wild success would look like being able to go through many important fields of science, find all the FRO-shaped bottlenecks, and solve a decent fraction of them.

FROs aim to address neglected bottlenecks in crucial fields like biomedicine and climate technology. Can you share a few specific examples of how FROs have successfully tackled these challenges and the impact they have had so far?

The field of bio-manufacturing aims to use biological organisms to replace unsustainable chemical processes with cleaner and less resource intensive ones powered by biology. However, most of the world's micro-organisms cannot be applied to this domain because we don't know how to grow them in industrial conditions outside of their natural habitats. The FRO, Cultivarium has developed a software for predicting key aspects of the recipes needed to grow diverse microorganisms in the lab directly from their genome sequence, which should accelerate finding growth conditions for many new organisms for bio-manufacturing.

If we could use artificial intelligence to do mathematical problem solving, we could accelerate many fields of basic and applied science and engineering such as developing new materials or energy sources. Yet state of the art AI still struggles with mathematical reasoning. Lean FRO has developed a software system for expressing and verifying math in code. This was recently used by researchers at the company DeepMind to bootstrap an AI program that is competitive on the international mathematical Olympiad.

There are many neuropsychiatric disorders which involve many brain areas communicating. Right now, we don't have a technology that can safely look at the dynamics of the whole human brain at reasonably high spatial resolution in real time during a patient's day to day life and use that to target tailored stimulation to multiple brain areas. That's what Forest Neurotech is aiming to build based on ultrasound on a chip technology. They've accomplished in less than a year what it took several BCI companies many years to achieve. This has been facilitated by the FRO structure which has enabled both academic and corporate collaborations that feed into this workstream.

You mentioned your interest in carbon capture and the need for more policy work to develop advanced markets and encourage company formation in this space. What are your thoughts on the most important policy changes needed to encourage market commitments for carbon capture within an economic development model?

One of our FROs, [C]worthy, is focused on the measurement, reporting and verification capability for ocean-based carbon removal. Without the ability to quantify the amount of carbon removal, and other effects, resulting from a given intervention in the ocean, there can be no market for such interventions, even if the government is the main payer. This inherently involves creating a complex piece of software, and building a community around that software, as this quantification must rely on modeling not just direct observation. So that's a technical change, but it requires a social ecosystem to be built around it, including policymakers. [C]worthy is focused on enabling this and making the technical work they do inclusive of and accessible to all these processes.

We're also excited about other recent efforts to connect the dots between fields to help regulators understand how to shape policy here in specific areas of carbon removal, because there are lots of details to get right and that may be method specific. Here is one such effort: https://www.carbonremovalstandards.org/about

*A Cell Observatory to reveal the subcellular foundations of life | Nature Methods