AUGUST 2022

BRIEFING PAPER

A NEW MODEL FOR SCIENCE

ten THE ENTREPRENEURS NETWORK

TONY BLAIR INSTITUTE FOR GLOBAL CHANGE

Convergent Research
A member of the Schmidt Futures Network

EXECUTIVE SUMMARY

MILAN CVITKOVIC, ADAM MARBLESTONE, AND ANASTASIA GAMICK
Convergent Research
HENRY FINGERHUT AND BENEDICT MACON-COONEY
Tony Blair Institute for Global Change
SAM DUMITRIU AND ANTON HOWES
The Entrepreneurs Network

– New types of scientific institution are needed to accelerate and unlock research progress and enable startups in critical technical fields like energy, medicine, and biosecurity.

– Focused Research Organisations (FROs) are an emerging model for transformative science projects. FROs undertake projects that are:

  – **too big** for a single academic lab to do,

  – **too complex** for a loose, multi-lab collaboration,

  – and **not directly profitable** enough for a venture-backed startup or industrial R&D project.

– Key features of FROs are:

  – They are run by full-time technical founders who oversee 10-30 employees

  – They pursue specific, quantifiable technical milestones rather than doing blue-sky research

  – They are finite-duration (5-7 years) efforts

  – As they near completion, they translate what they have built into venture-backed startup spinouts and/or longer-lived nonprofits

– The UK has unique opportunities to leverage the FRO model in biomedicine and net-zero carbon technologies. This is thanks to research infrastructure like the UK Biobank and NHS data sharing programmes; ambitious, innovation-oriented programmes such as the Net Zero plan and ARIA; and research translation and technology dissemination mechanisms such as the Catapult Network and NHS Transformation Directorate.

– ARIA and other funding bodies in the UK should consider partnering with existing philanthropic efforts and launching FROs in biomedicine and net-zero carbon technologies.

Introduction

Academic research groups and startup companies are the workhorses of translational science, bringing technological innovations out of the lab and into our lives. But there are some loads you cannot ask a workhorse to pull. A university astronomy lab could not have launched the Hubble telescope, nor could a venture-backed startup have built the Large Hadron Collider at CERN.

Today, scores of worthwhile scientific endeavours remain unattempted because they do not fit naturally within existing scientific institutions. In fields critical to our health and wealth in the coming decades—fields like green energy, medicine, neurotechnology, and biosecurity—society has mostly given up on institutional innovation in science, seemingly expecting miracles from the status quo.

The UK Biobank is an instructive exception. Despite facing initial criticism for departing from "standard academic scientific practice" in its organisational structure, the Biobank is among the most influential biomedical initiatives in world history.[1] Over the past decade, 28,000 approved researchers from 86 countries have used the Biobank to publish 4,600 papers and create countless new therapeutics and biotech startups. It is neither a university lab project nor a startup company, and could not have been done as either. Rather, it is an independent PLC and registered charity that executes like a startup, led by a technical CEO who also happens to be a Professor of Medicine and Epidemiology.

The UK Biobank highlights a common pattern in translational science: a need for projects that are **too big** for a single academic lab to do, **too complex** for a loose, multi-lab collaboration, and **not directly profitable** enough for a venture-backed startup or industrial R&D project.

If the UK wants to achieve its potential as a leader in scientific and technological innovation, it needs to embrace institutional innovation. It has already made a start with initiatives like founding the Advanced Research and Invention Agency (ARIA), but it can go even further. Specifically, **Focused Research Organisations (FROs)**[2] are a new type of scientific institution designed to fill the gap between academic science and venture-backed startups. In what follows, we argue that FROs should be part of the UK's ambitious scientific and technology push.

1 A contemporary history of the origins and development of UK Biobank 1998-2005, Mairi Anna Langan

2 Unblock research bottlenecks with non-profit start-ups, Nature Comment, 2022

Focused Research Organisations: An Opportunity for Institutional Innovation

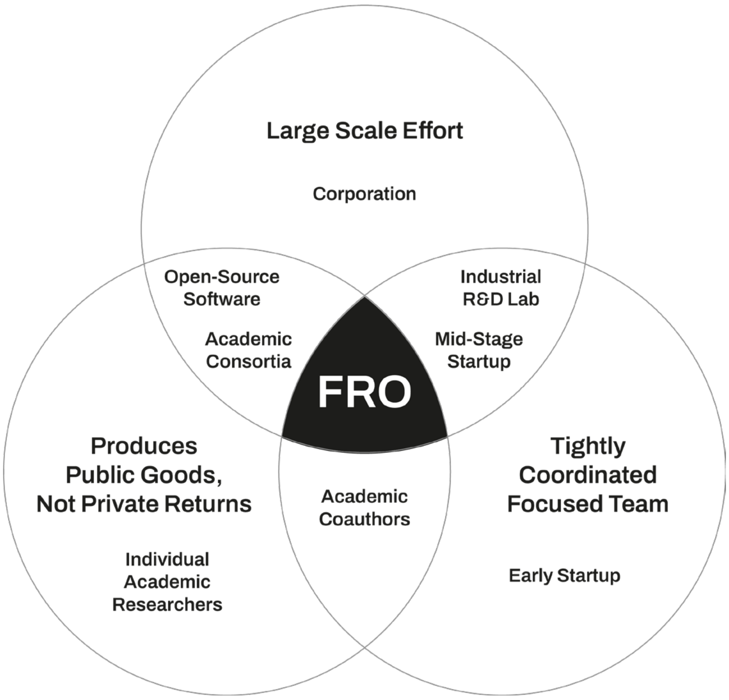

FROs are essentially nonprofit tech startups. As shown in the diagram, they fill a gap in the institutional landscape of science: a missing type of organisation for **large-scale, tightly coordinated, non-profit projects.**

While it is true that National Labs or institutes sometimes pursue large, focused initiatives internally, the FRO model allows anyone with a good idea and team to do so. Just as startups make it possible to pursue ambitious business ideas outside of large companies, FROs democratise and catalyse the pursuit of ambitious science projects.

In brief, the defining features of an FRO are that:

– They are led by a full-time **founding team** of scientists, engineers, and entrepreneurs.

– They consist of a **larger-than-academic-scale** team of 10 to 30 (or even more) interdisciplinary scientists, engineers, and project managers.

– They are "focused" in that they have **specific, quantifiable technical milestones** they must achieve rather than doing blue-sky research.

– They produce **high-impact public goods** for science and technology—massive datasets, next-generation research hardware, open-source production protocols, etc.—rather than trying to capture value from a marketable product

– They are **finite-duration** (5-7 year) efforts to avoid mission creep and preserve focus
– As they near completion, they **actively translate** what they have built into one or more venture-backed startup spinouts and/or longer-lived nonprofits.

Leadership by a full-time founding team is important. FROs are designed to do technically challenging work comparable to that of an early tech startup, and founder-led startups have historically outperformed others.[3] Indeed, if a project does not require the active management and singularity of focus that a founder-led startup does, it probably doesn't need to be done as an FRO.

It should be emphasised that entrepreneurs, managers, and operators are as critical a part of FRO founding teams as scientists. The coordination and execution intensity required to run an FRO exceeds business-as-usual science. In a standard academic lab, the principal investigator mentors students on individual or small group projects. In contrast, FROs have teams of 10 to 30 (in some cases, more) employees with diverse technical skill sets, requiring tight coordination to achieve their goals. Managing such a team requires the skills of startup entrepreneurs.

Another similarity to startups, denoted by the "Focused" in "Focused Research Organisation," is that FROs are not intended for open-ended scientific exploration. FROs specify quantifiable technical milestones at their outset, and the FRO's success or failure — and possibly continued funding — is determined based on meeting them. While an FRO may pivot its strategy for achieving its goals, the goals do not change. Having fixed goals is important for team focus and for making the projects legible to funders. None of this is to impugn open-ended inquiry, which is the bedrock on which science is built. FROs simply fill a vacant niche in the scientific landscape, whereas open-ended research is well supported in existing academic structures.

One difference between FROs and startups is their nonprofit status. If a hard technical R&D project has a path to profitability, it's better for it to be supported by the market than to rely on philanthropy. But not all projects can capture enough of the value they create to become profitable. Examples from recent scientific history include the Neuropixels probe in neuroscience, the Tensorflow open-source library for machine learning, or the UK Biobank dataset mentioned earlier. Thousands of labs use these tools every day, and dozens of companies have used them to build medical devices, AI systems, and therapeutics, respectively. But each cost far more to develop than could ever be recouped by selling them to academic labs and early-stage startups.

Another difference between FROs and startups is their finite duration, which is specified along with its technical milestones. Typically this is around 5 years. FROs' fixed durations make them clear, bounded requests

3 Founder CEOs and Innovation: Evidence from S&P 500 Firms.

of funders. FRO funders know they won't be pestered for additional capital indefinitely and can worry less about mission creep and institutional rot diminishing the organisation's impact. An FRO's fixed duration (and mostly upfront funding) also means FRO founders don't have to spend time continually fundraising or adapting the mission to justify the continued existence of the nonprofit. It also differentiates them from National Labs and institutes like Germany's Max Planck Institute and the UK's Turing Institute, which provide infrastructure and user facilities for longer-lived projects than an FRO would undertake.

Every FRO will have a different plan for how to maximise the impact of what they produce after the conclusion of the project. Some FROs, like ones that produce large datasets or experimental protocols, will likely have their highest impact by openly sharing this data. Such FROs may hand over the fruits of their efforts to existing institutions for stewardship, or start longer-lived, smaller nonprofits for this purpose. Other FROs may produce Intellectual Property that will best serve the organisation's goals by being licensed to one or more startups. In such cases, founders and employees of an FRO would be ideal founders of spinout startups intended to develop the technology unlocked by the FRO.

Extant FROs

As of writing, three FROs have been launched, each funded in the tens of millions USD, all in the United States. Thus far all FROs have been funded by private philanthropy, but they can just as easily be funded by a governmental entity like UKRI or through a public-private partnership.

E11 BIO

- Mission: build the key tools needed to map the connections between every neuron in a mammalian brain
- Impact if successful: new treatments for brain disorders, new experimental paradigms in neuroscience, new applications in brain-inspired computing
- Duration: 5 years
- Location: SF Bay Area, USA

CULTIVARIUM

- Mission: build an end-to-end toolkit for cultivating currently-unculturable microbes
- Impact if successful: accelerate the study and engineering of microorganisms for applications in medicine, carbon removal, and beyond
- Duration: 5 years
- Location: Boston, USA

REJUVENOME

- Mission: conduct the largest study of ageing in animals ever performed
- Impact if successful: provide the field of ageing research a gold-standard dataset on which to base future work
- Duration: 7 years
- Location: SF Bay Area, USA

Focused Research Organisations in the UK

The UK's strong scientific and entrepreneurial talent base position it to take full advantage of the FRO model, with unique opportunities to specialise in biomedicine and net-zero carbon technologies.

Biomedicine

PARTNERING WITH THE NHS

The potential of the NHS for biomedical research is extraordinary and unique to the UK. Its size, centralised approval processes, and consistency in data and standard of care make it possible to run biomedical projects in the UK that could not be run elsewhere.

Metagenomic Sequencing

For example, the NHS would also be an ideal partner for developing novel pandemic prevention and biosecurity technologies. One such technology is "metagenomic sequencing", in which waterways and wastewater are ubiquitously and continuously sampled and watched for any exponentially spreading DNA or RNA sequence, the presence of which would indicate an emerging biological threat.

No system for automated sampling, sample preparation, and sequencing like this exists. An FRO would be ideally suited to design one because its development is too engineering-heavy for university labs to perform, but too early-stage to be commercially viable.

Such an FRO would need to be tightly integrated with the NHS, because to test whether it is successfully spotting diseases as they spread across the UK, the FRO would need to compare its findings to real-time health and testing data from hospitals nationwide. This would also be an opportunity for the NHS to test and refine the data sharing infrastructure it is building under the General Practice Data for Planning and Research programme. Additional partnerships could be made with key UK air and sea ports to track disease strains as they enter the country.

Pandemic Practice Runs

Another timely use of the FRO model would be performing pandemic "practice runs": trying to set new speed records in vaccine or drug development, manufacturing, and distribution. Such a project has no immediate commercial value and falls outside the purview of academia. An

FRO could partner with the NHS and other UK scientific resources like the Diamond light source, which acted quickly to provide useful preclinical data in the beginning of the COVID-19 pandemic. The UK's leading Vaccine Manufacturing and Innovation Centre would also be a strong partner for this FRO.

Wearable Laboratories for Public Health

In terms of more general public health risks, an FRO might be established to build technologies for diagnostics and biomonitoring. Technologies like the continuous glucose monitors worn by diabetics or smart watches have the potential to monitor hundreds or thousands of biomarkers rather than the handful they can now. But most diagnostic and monitoring technologies languish in academic labs due to a chicken-and-egg problem: they aren't commercializable until a clinical benefit is demonstrated, but to find a clinical benefit one has to manufacture the technology at commercial scale and quality in order to run a large clinical study. An FRO could break this cycle by engineering and manufacturing new monitoring devices at scale for research purposes.

Where does the NHS come in? Discovering the best use-cases for these new monitoring devices will require running prospective clinical trials to connect the data obtained by the devices with as many health outcomes and other biomarkers as possible. These trials will need to be large-scale, have diverse cohorts, and, critically, have consistent standards of care between the trial participants. Running such a trial would be almost unthinkable in a fractured healthcare system like that in the United States. Not so in the NHS, as demonstrated in the recent RECOVERY Trial. And digital health initiatives across the NHS at the Trust-level, region-level (e.g. DigitalHealth.London), and system-level (e.g. NHSX, now part of the Transformation Directorate) provide an opportunity to test wearable technologies under real-world conditions.

UK BIOBANK

As mentioned above, the UK Biobank houses a dataset without equal in biomedicine: health histories and genetic data from almost 500,000 individuals from across the UK. But unlocking the potential of the Biobank requires not just hosting data, but algorithm development and big-data engineering that are beyond the capabilities of the academic labs who use it. There are a number of potential FROs that might develop it further.

Biobank Family Study

One example would be to augment the dataset by adding first-degree relatives of Biobank participants: parents, full siblings, or children, with whom they share 50% of their genes on average. The value of this inclusion would be that studying groups of first-degree relatives reduces the chances of discovering spurious correlations.

For example, if a researcher were to notice that people in the Biobank with a particular genetic mutation have a high propensity to develop breast cancer, it could be that the mutation changes their biology in a way that causes breast cancer. But it could also just be that a particular socioeconomic group or geographic area has a high occurrence of this mutation, and *coincidentally* also has a high propensity for breast cancer for some other reason not related to the mutation. These confounding biases have been shown to lead to mistaken inferences of the causes of diseases and have misled researchers for decades.

Studying first-degree relatives in genomics research helps avoid spurious correlations because the health outcomes of full siblings in a family are like a mini randomised controlled trial of the effects of the genes of their parents. When two full siblings differ in a health outcome, and only one inherited a particular mutation from their parents, it is much stronger evidence that the mutation played a causal role in the health outcome than if the two were unrelated strangers.

The causal understanding gleaned from a first-degree relative study at the scale of the UK Biobank would revolutionise genetic science, just as randomised controlled trials revolutionised medicine and epidemiology. But collecting genetic and health data on first-degree relatives of Biobank participants would require both funding and personnel beyond that available to an academic lab, as would the software engineering resources needed to analyse such a large amount of data. An FRO, however, could partner with the UK Biobank to perform a first-degree relative dataset augmentation and build the high-performance computational tools necessary to analyse it too. Partnerships with the National Institute for Health Research Applied Research Collaborations or its close collaboration with Genomics England could also be fruitful.

Climate Technology

The development of climate technologies was one of the original motivations for the FRO model, because their timelines are so much longer than those used by venture capital. A standard VC fund operates on a 10-year timeline in which they need to invest all of their capital and then return the profits to the fund's investors. But testing an intervention like olivine weathering or fusion power could take decades.

Oceanic Carbon Sequestration

Given its maritime capability and varied territorial waters, the UK would be an ideal jurisdiction for FROs studying oceanic carbon sequestration strategies – that is, ways to store carbon dioxide in the sea. Oceanic sequestration strategies are among the most promising routes to achieving the levels of CO2 removal from the atmosphere that the UK is legally committed to reaching by 2050 per its Net Zero plan and COP26 pledges. Many strategies have been suggested, including stimulating growth of algae or enhancing the rate of biomass sinking with clay minerals.

As promising as ocean sequestration is, it carries significant "unknown unknown" risks, again making an FRO better suited to the task than a startup: thorough experimentation well beyond the scale of academia is needed to ensure its safety and efficacy. Such experimentation would be greatly improved by testing in diverse marine ecosystems, and the UK's waters encompass as many marine ecosystems as any nation.

Geothermal Everywhere

The UK also has great potential for developing deep geothermal energy technologies. Though often overlooked in the green energy discussion, there is 23,800 times as much geothermal energy in the Earth's crust as there is chemical energy in all the Earth's fossil fuels. Existing geothermal power plants require hot water to be located relatively shallow in the earth, as it is in Cornwall where the United Downs Deep Geothermal Power plant is located. But a number of drilling technologies have been proposed that would allow digging deep enough that nearly any location in the UK would be a viable location for a geothermal power plant. These technologies could eventually power the entire UK with safe, zero-carbon, low-footprint, 24-hour baseload power. But the relatively low returns of the energy industry make investing in large-scale drilling R&D infeasible. FROs in next-generation geothermal energy could take on this research and engineering task, unlocking a geothermal revolution in the UK — a fitting legacy for the country which invented deep mining in the first place.

Clean Meat

On the agriculture front, the UK's food system accounts for 30% of domestic emissions, and meat has the highest greenhouse gas emissions per gram of protein. Cultured meat technologies could drastically reduce these emissions, but significant technical challenges remain. One challenge that an FRO could help overcome is developing open-access formulations for the feedstocks that cells subsist on as they grow into meat. Existing formulations are underdeveloped and underprovided by the market given that cultured meat remains pre-commercial. But climate change will not wait, and the availability of more open-access formulations—especially those that have been optimised for species and cell types relevant to cultivated meat—will provide a foundation to enable both academic researchers and startup companies to develop their own customised formulations with far less effort and cost.

Conclusion

Focused Research Organisations give UK entrepreneurs, scientists, and engineers a new path for developing transformative technologies, one that fills a vacant niche in the existing science and technology translation ecosystem. The FRO model can be applied across scientific or technological fields, with funding from single sources, consortia, or government, with agile teams and minimal overhead.

If aspiring founders of a technical or scientific project feel like they are deforming or curtailing their idea to fit it into the startup or academic mould, they should consider executing their project as an FRO. This is doubly true if the project is in biomedicine or climate technology, areas where the UK has unique advantages for pursuing FROs. Interested potential founders should contact Convergent Research,[4] a member of the Schmidt Futures Network. Convergent Research is a nonprofit dedicated to incubating, launching, and supporting FROs.

As important as the founders is the funding. Several of the projects described above are shovel-ready, with teams ready to launch their FROs the moment funding is secured. But they need funders who are willing to experiment with new organisational models for science.

The UK Government should take a leading role here. They have recognised that the existing process of funding science isn't working as well as it could, taking steps such as launching a review into Scientific Bureaucracy and founding ARIA. Such initiatives are laudable, but they alone are not a complete solution to the structural barriers that plague scientific research and early-stage technology development. To this end, ARIA or UKRI should experiment with deploying their capital in the form of FROs. Where better than in the home of empiricism to experiment with how science is organised?

4 Convergent Research. convergentresearch.org

Convergent Research

A member of the Schmidt Futures Network

TONY BLAIR INSTITUTE FOR GLOBAL CHANGE

ten THE ENTREPRENEURS NETWORK

The Entrepreneurs Network
84 Eccleston Square
Pimlico
London
SW1V 1PX