

High Quality Research Environments

Professor John Gallacher,
Director, Dementias Platform UK
University of Oxford

Professor Chris Webster,
Dean of Architecture
Hong Kong University

Media Summary

High quality research environments are critical to scientific progress and maximising public benefit. This paper argues that the increasing complexity of emerging questions requires more efficient collaborative structures. Economic club theory is used to identify ways of incentivising collaboration including increased data access, trust-by-design solutions, and user-led innovation. It is acknowledged that good science is based on both aesthetic and ethical judgements and that environments that combine these are likely to produce motivated scientists and impactful research

Abstract

A major challenge facing all research communities is creating and sustaining high quality research environments. A model describing strategic social structures that constrain knowledge production suggests that targeting these structures will have greater impact than addressing issues surrounding individual lab cultures, as important as these are.

A literature search identified five common themes underlying bioscience research environments comprising collaboration, data processing, confidence in data and scientists, trust, user-led development, and a deep commitment to public benefit. Club theory was used to develop a model describing the social structures that constrain and contextualise research environments.

It is argued that collaboration underlies impactful science and that this is hindered by high transaction costs, and the benefits associated with competition. These combined with poorly defined property rights surrounding publicly funded data limit the ability of data markets to operate efficiently. Although the science community is best placed to provide solutions for these issues, incentivisation by funding agencies to increase the benefits of collaboration will be an accelerator. Given the complexity of emerging datasets and the collaborations need to exploit them, trust-by-design solutions are suggested. The underlying 'glue' that holds this activity together is the aesthetic and ethical value-base underlying good science.

The Collaborative imperative

Science illustrates the compelling advantage of collaboration whether at research team or consortium level. Every further step in specialisation is a step towards greater interdependence as scientists with specialist knowledge need scientists with other specialist knowledge to address emerging questions. And herein lies a paradox. Or at least a challenge. Access to limited resources like a faculty position, research funds or a high-impact publication drive competition, but, *a priori*, because knowledge is fragmented, collaboration between individuals and groups will improve the final product. The extent to which competition *per se* has generated insights that would not otherwise have been made, is unknowable. In contrast, the extent to which collaboration has generated insights that would not otherwise be made is easily measured by the author list on any peer review publication. These lists measure the revealed preference for collaboration over exclusivity in creating complex knowledge. Author lists reveal how the competition-collaboration paradox works itself out in practice. Optimising growth of a data economy rests, therefore, on optimising the balance between individual and collective incentivisation.

In the absence of any meaningful data, we explore this problem in a thought experiment that models a scientific collaboration as a club⁶, where decisions are made by individuals regarding the value of joining the club (joining a collaboration), and by club members on the value of extending membership (increasing the number of collaborators). In this model the benefit of collaboration is the ability to address specific scientific questions with greater precision at lower cost.

Figure 1 shows how net benefits and overall costs; production and transaction costs (vertical axis) vary according to complexity; the number of collaborators (horizontal axis). Initially, costs (green solid curve), largely driven by the costs of knowledge production, reduce with collaboration through the wider pooling of specialised knowledge and a more distributed workload. However, collaboration itself is not cost free. As the club increases in complexity (size), the overall costs increase, due to the growing costs of transaction (distributing responsibilities and rewards, ceding peer-esteem indicators to others, etc.). This change in costs affects net benefits (red solid curve). At some point net benefits will be outweighed by transaction costs and collaboration will stall. The intersection of net benefit and overall cost curves identifies the upper collaborative limit (blue solid line). Figure 1 also shows where the difference between net benefits and overall costs is greatest, suggesting the point of optimal collaboration (grey solid line).

Two parallel strategies can be used to extend both the upper translation limit (blue dashed line) and the optimal point of collaboration (grey dashed line). Infrastructure can be used to reduce transaction costs (red dashed curve). *A priori*, a good platform will increase the average size of cooperating groups and is likely to yield more scientific product per data asset. UK Biobank provides an example, reducing search, provenance, legal and many other costs of accessing complex, high quality longitudinal data.

As a result, this platform has become one of the world's most used biomedical resources. A second strategy is to engineer more socially optimal decision-making into the value chain (green dashed curve). For

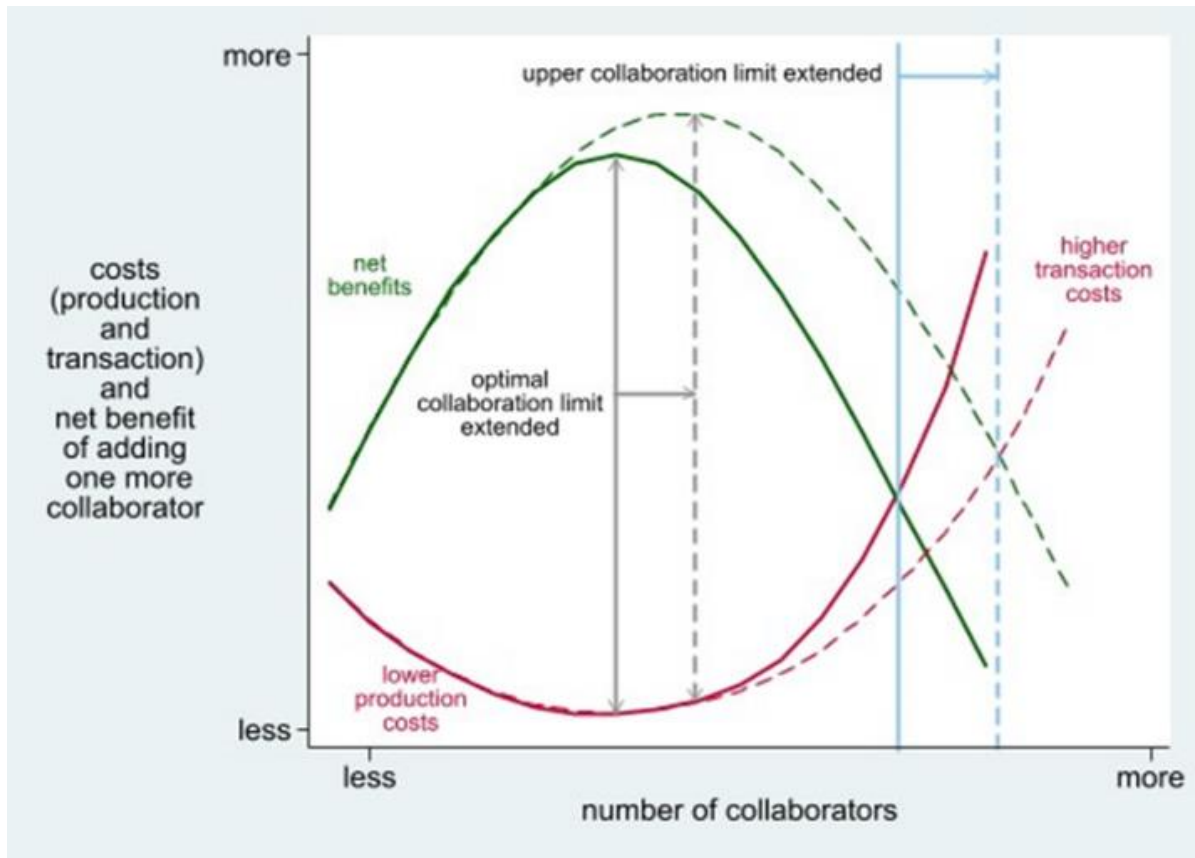


Figure 1:
Schematic of extending the collaborative limit by reducing transaction costs
and increasing collaborative incentives
(adapted from James Buchanan [6] and Richard Cornes and Todd Sandler [10])

example, preferring proposals that use existing infrastructure, and using metrics of third-party data usage to evaluate overall scientific impact. Of interest, both strategies facilitate spontaneous decentralised collaborative solutions. Also of interest, the model provides a framework for understanding the optimum size and organisation of infrastructure.

Current incentive structures are strongly influenced by what may be described as the competitive dividend. This is not necessarily helpful. In a review of Australian biomedical grant proposals, Herbert et al. report that the opportunity cost for preparing new proposals averages at 38 working days, with a 79% failure

rate⁷. For the 2,966 failed proposals identified in the study, this translated into an annual salary cost of AU\$52m. Arguably, structures that emphasise the collaborative dividend, resulting in fewer, higher quality proposals, are likely to deliver greater overall benefit at reduced cost.

These arguments illustrate that collaboration decisions are typically made at the margins, i.e. in terms of benefit to individual scientists, rather than benefit to the wider scientific community. The extent to which a marginal decision-making calculus will differ from an overall decision-making calculus will vary across scientists, but they are not identical. Modelling the impact of collaboration on specific private

and social benefits can be used to inform strategic funding decisions on consortium formation and the organisation of infrastructure.

Making data markets work

In demonstrating the logic of private as well as social benefit in the development of consortia, Figure 1 also describes data-sharing behaviour in practice. The largest barrier to scientific advance may be considered to be the failure of the data market to provide access to those who can add value to data already collected. In a survey of 3,556 articles from 333 open access journals Gabelica et al⁸ found that only 7% of corresponding authors responded positively to a data access request, even when their intention to share data was explicitly stated, meaning that in this experiment, the revealed preference of 93% of authors was not to share even if their stated preference was to share.

If clubs of scientists act to preserve internal shared value of data, it is of interest to consider the mechanisms they use to do so. Access to products in data markets is controlled by informal and formal rules that can be thought of as allocating property rights (a right to benefit from a resource). Clearly defined property rights reduce the costs of transacting in markets by providing a secure basis for decision-making and thus incentivise creativity and innovation. But extant institutions do not necessarily allocate property rights to scientific data efficiently, with respect to either private or social benefit.

Assigning property rights to data is a social preference. There is no correct approach. But some preferences are more socially efficient than others in terms of maximising knowledge creation, whilst preserving the interests of the various

stakeholders. A rule for socially efficient knowledge creation would require rights to be assigned to those in the best position to use the resource for a desired outcome, such as 'creating scientific opportunity', or 'maximising the number of published studies', or 'maximising the chance of translating discoveries into impacts. Clearly such a rule is not typically in place currently, with *de facto* rights to further exploit data typically being held by the club of scientists who first created or came to hold ownership of a dataset. In this important sense, for publicly funded science, property rights are unclear⁹. For example, although a dataset is legally 'owned' by a host institution, it is considered by funders (and the general public) as held in trust for public benefit. Problematically, as Gabelica et al⁹ showed, what constitutes 'public benefit' is capable of broad interpretation.

There may be ambiguity in ownership of several kinds of rights: a funder's right to require public trust responsibility; a university's right to uphold this; a researcher's IPR over work that adds value to data; a researcher's right to subsequent or subsidiary research using the same data; a researcher's public trust responsibilities; and so on. All of these are important and it is worth considering how they might be best allocated to different stakeholders to support the growth of the data market.

Figure 2 categorises data on two dimensions that relate to the manner in which data are consumed by users: a) rivalrous consumption (congestible or depletable) and b) excludability (technical, legal, or political). A 'club', as an organisational form, is an efficient mechanism for organising goods that are non-rivalrous and excludable (cell 2)¹⁰. By excluding non-members, data can be used without congestion and without rivalry within the club. Publicly funded data which

	Excludable	Non-excludable
Rivalrous (depletable)	1. Private Good (privately funded science resource)	3. Common pool resource
Non-rivalrous (Non-depletable)	2. Club Good (publicly funded data and biosamples)	4. Public Good (astrophysics data)

Figure 2:

The data market

though not depletable, is congestible in the sense that it might be subject to mishandling to public harm, threatening continued existence or use of the data. Club ownership is an efficient solution in such circumstances, with club rights allocated by scientific competence and *bona fide* purpose. Sample collections could be considered rivalrous and depletable (cell 3) with open access. However, with limited access, they become more sustainable, giving further justification for club ownership, and requiring perhaps stricter membership conditions.

However, data that are technically excludable but unlikely to be congested by open use, are inefficiently organised by either scientific clubs or privately by the data's creator. As a general strategy, restricting property rights to such data is a low-growth, low innovation path. In the late 1980s, the UK Ordnance Survey took proprietary ownership of its national spatial databases and charged eye-watering sums to academic, local government and industry users. The USA took the other approach, making a raft of location-based databases free to all users who would add value. The Geographical Data Science community and industry in the USA consequently exploded in growth that was several orders of magnitude

greater than in the UK. Non-congestible scientific data that are technically excludable are best not made exclusive. Changes in data access rights, platform management protocols, grant condition and so on can be better designed to remove exclusive access after the period during which exclusive access is appropriate. Within science communities there is a lack of clarity in switching rights from exclusive to the public domain. This arises because funders with a fiduciary responsibility to maximise social benefit from research funds give away too much of their right to grant bidders, who capture it for their own benefit.

We should also note, however, that some data are both non-rivalrous and non-excludable (cell 4). These are pure public goods in terms of consumption characteristic and are most efficiently placed in the public domain, i.e. with rights given to all. Some data become like this over time, where the value added to data by multiple, entrepreneurial, scientists and citizens does not devalue the data for others. In this, astrophysics has an honourable history.

At the other extreme, some data are most efficiently organised as a private good when access can be technically, legally, and politically, denied (cell 1). Typically, these

data have high novelty, sensitivity and commercial value, and are expensive to create. Private ownership is crucial to maintaining the production of such data until such time as the costs of creation has been recovered, or reasonable 'creator rights' have been exhausted, or the data replicated more cheaply by new technology.

These distinctions allow the application of non-arbitrary and transparent rules of access. Although assigning property rights to those best placed to achieve the desired outcome are user-community judgements, the principles on which these judgments can be made more explicit.

While clubs are the natural order in the scientific data economy, they are not necessarily the most efficient mechanism for allocating rights to every kind of data. Appropriately adjusting access, use and ownership rights can be addressed implicitly and explicitly. Dementias Platform UK (DPUK) addresses it implicitly by reducing the reputational and other risks associated with data access; leading to an access request approval rate (after triage for egregious proposals) of 90%. UK Biobank addresses it explicitly through a standard data controllership agreement; leading to an approval rate of 99%. These platforms effectively manage responsible access on behalf of multiple contributors, achieving economies of scale and protecting against data 'congestion'. Platforms can be thought of as shifting data along the continuum from restrictive club good, to less restrictive club good via a larger club and federations of clubs to achieve greater social benefit. Where a platform triage and approval \approx 100% of applications, it has effectively converted an inefficient club good into a public good.

Major shifts in systems of rights that cohere society often follow periods of great upheaval. Prior to SARS-CoV-2,

linking UK administrative health data to research data required months of negotiation involving serial licence fees. However, through the pandemic, use rights and procedures were rapidly reviewed, enabling accelerated access resulting in globally impactful science¹¹. SARS-CoV-2 has demonstrated the high value of making these data widely available, and the ability to do so at pace. However, accessing these data for non-Covid purposes remains tortuous and expensive, even though other conditions such as dementia, a slowly unfolding global pandemic, are much greater and enduring public health challenges. This serves to highlight data-pipeline inefficiencies in the sense we have been discussing. It may be that the social conversation about improving resilience to future pandemics, tips the scales in favour of a more clearly defined approach to data access.

Constrained user-led innovation

The organisation of scientists into clubs of different size and function evolves over time. Never at equilibrium, clubs will vary in size, longevity, and configuration. As we have discussed these trajectories will be determined by the dynamics of benefits and costs. For those concerned with building high quality research environments, a strategic question is can we grow the science-base through manipulating the organisation of scientific production?

'Top-down' solutions typically lack agility and responsiveness in the face of changing needs, whilst spontaneous, decentralised, user-led solutions are hallmarks of an innovative research culture. However, in a data-driven economy, uncoordinated and competitive user-led solutions frequently generate perverse outcomes. These include multiple idiosyncratic data models, poorly annotated data, poorly documented metadata, and duplicative

non-standard data processing pipelines. All of these contribute to increased production and transaction costs through extended research cycles and non-reproducible analyses. In this way, too much decentralised control makes accessing the data market more difficult. This is a dilemma.

To address this, it is helpful to understand that successful decentralised systems of exchange require well-oiled sets of enabling institutions that create an ecosystem to underpin efficient collective behaviour. From informal rules that operate within individual labs, to legal obligations regarding data access, agreements emerge that constrain some activities in order to promote others. These institutions form the structure that governs scientific activity. In a data driven economy, institutions governing data standards deserve close attention. In market-driven sectors this tends quickly to lead to industry-organised standards. Without standards, for example, regarding the colour of 'live' wires in electrical goods, spontaneous exchange of data and goods rapidly slows.

In the biosciences, an example is reference SNP cluster ID (rs) numbers¹². By establishing rules around how (and how not) to annotate genetic data, confidence in the provenance of data is increased, transaction costs of data access are lowered, and the pooling of property rights to support rapid publication is incentivised. The introduction of rs numbers has been transformative. From this example, two lessons can be learned. The first is that institutions to support scientific activity emerge according to need; Rs numbers are a response to the need for greater reproducibility. Second, the emergence of institutions can be accelerated by the leadership of stakeholders with critical mass. Interestingly, although in the case of

rs numbers, the scientific community generated the solution, it was coordinated action by two major stakeholders (National Center for Biotechnology Information and National Human Genome Research Institute) that realised the ambition. Without this leadership, it is likely that rs numbers would have been just one of a number of (equally useful) solutions, each with its own band of disciples.

Arguably, non-proprietary research funders exist to support collective scientific goals and activity, and have a role in supporting the science community to develop efficient institutions. Not least because these make science attractive to potential scientists, avoid wasting resources, and grow the national innovation base. However, funders tend to eschew this subtle but critical role, as it can be interpreted as an imposition on academic freedoms. An alternative perspective is that funders have a responsibility to work with scientists to support the development of mature institutions that provide an efficient science data ecosystem.

Consider the example of research cohort data. Cohort studies are a critical component of biomedical science's armamentarium. The growing importance of cohorts is demonstrated in that, for dementia alone, the number of cohort-based publications per year increases monotonically; currently exceeding 2,000¹³. Typically, each cohort uses a bespoke data model and governance structure that has evolved over time according to its scientific priorities and resource constraints. This 'exercise in entropy' incurs substantial knowledge production costs and leads to significant transaction costs for third-party researchers. Although the value of curating cohort data to a defined standard is widely accepted, few research groups have been willing to invest

in this^{14,15}. The rise of multi-cohort data management platforms such as DPUK^{16,17}, Dementias Platform Australia,¹⁸ and the Alzheimer's Disease Data Initiative (ADDI) Workbench,¹⁹ has demonstrably improved the 'public good' quality of data created by federating clubs of scientists; and provide a model for improvement. The argument is not that some higher authority should coerce data providers to invest in something that is not in their interest. It is that funders can steer scientists to engage in win-win collaboration through lowering systemic transaction costs of data exchange and increase the public good value of data they fund.

Trust-by-design

Trust is the implicit operating principle underlying human collective activity. Trust simplifies otherwise complex and unpredictable environments; identifying points of certainty around which to organise; and agreeing on a culture where key uncertainties are removed on the basis of mutual agreement, respect and ethical codes. By facilitating better prediction of the likely reciprocal behaviour of others in sharing costs and benefits, trust fosters collaboration, and the efficient assignment of informal and formal property rights. Trust is also foundational to the provenance of data and technologies. As datasets grow in size, complexity and sensitivity, confidence in the provenance chain becomes increasingly central to the viability of the entire science economy.

Informal trust-based solutions work well for bilateral collaboration and can work for small communities. However, emerging research questions frequently require multi-lateral collaboration involving large numbers of diverse stakeholders, and for these, time-tested less formal solutions are inadequate. Multiple actors with multiple interests involving multiple data-sources generate complexity leading to potentially

prohibitive transaction costs. An alternative strategy for managing multi-lateral collaboration is trust-by-design. Here, legal, privacy, security and scientific requirements are embedded within technical and organisational workflows, that are explicit, transparent, and fully auditable. This enables systematic streamlining, standardisation, and automation. Although designed to service multi-lateral collaboration, embedded workflows have utility for collaboration in general. Trust-by-design solutions provide the information necessary for accurate and rapid judgements of trustworthiness and scientific value. It involves a shift from trust in a person or a group, to trust in a system. It is a process that underpins all but the most primitive of economies and is fundamental for human collective advancement in all senses.

For data management, trust-by-design solutions include trusted research environments (TREs) where data are stored and processed and accessed within a secure and fully auditable context. In the DPUK Data Portal, for example, data are ingested, curated, and made available to third-party researches upon approval from data controllers. Data cannot be downloaded and discovery, access, and analysis occur remotely. Standard legal agreements cover data deposit and data access, and access rights can be withdrawn if those agreements are abrogated¹⁵. Trust-by-design solutions also provide a level of familiarity and certainty that further extend trust. In DPUK, this enables rapid decision-making with a median decision time of 23 days. With the increasing complexity and sensitivity of emerging datasets, trust-by-design is likely to prove the only scientifically and socially acceptable solution for at-scale research environments.

Blockchain technology is a similar game-changing collaborative technology. In principle, it provides a technology for vastly reducing transaction costs of data access by embedding formerly lengthy administrative procedures into smart contracts. In distributing legal agreements, access rights, and transactional records across all users, it obviates the need for trust in a single organising authority. It is not yet a proven technology for scientific use, but there is no reason why that should not happen.

On values

Our brief survey of abstracts shows that to many scientists it is important that their work contributes some form of scientific and social benefit. It is worth considering what this does and does not mean for research environments.

The intellectual underpinning of science comprises observation, causal inference, and application. As scientists, we are familiar with the systematic observation of phenomena and consider this delivers objectivity (ontological order). We are also familiar with deriving insight by applying causal inference to observation (epistemological order). Both activities are considered intrinsically aesthetically satisfying. Furthermore, their utility lies in being untrammelled with extrinsic judgements of social benefit.

For application, however, the situation is reversed. Wittgenstein observed that if we compiled an exhaustive book of all the observations in the universe, and the insights derived, it would not contain a single objective value statement²⁰ i.e. not provide a basis for objectively evaluating social benefit. It appears, that as scientists, we are compelled to make judgements of value (regarding social benefit or otherwise) that cannot ultimately be derived from observation or causal

inference. Responses to this predicament vary, and their expression is contingent on the state of knowledge. Nevertheless, at the root of every value chain there lies what may be described as an ethical judgement, be it based on a belief in some ultimate value, or deontological framework, or some other more capricious basis of judgement.

Wittgenstein's exposition is extreme, as it does not acknowledge the rich tapestry of instrumental (utilitarian) values that we use to guide marginal decision-making. But, it helps make the point that for research environments, confusing aesthetic and ethical judgements is unhelpful. For example, the desire for inclusive science has both aesthetic benefit (more informative data) and ethical benefit (reduced inequality). They are closely entwined, but are distinct, involving different and sometimes conflicting priorities that need to be clearly understood for optimal decision making. Nevertheless, for research environments, systems and practices that address both are likely to be most attractive. The more deeply we understand the aesthetic and ethical implication of our research, the better positioned we are to develop a research environment that motivates individual scientists and clubs thereof, towards a higher calling.

But perhaps, for the consideration of research funders and others in a position to lead a change of culture, it more powerfully begins with the ethical. To acknowledge the dignity of the individual in terms of academic attribution, training, career structure, etc. will create a cadre of collaborative, efficient, innovative and trustworthy scientists who are better positioned to conduct rigorous science for social as well as private good. If a highlight of being a scientist is belonging to a truly international community that transcends

race and creed, then an intrinsic motivation exists to realise this ideal, in which aesthetic and ethical motivations coalesce.

Towards high quality environments

The science economy described here informs a strategic approach to developing high quality research environments. It identifies key socio-economic structures that contextualise and constrain the behaviour of research groups and individuals. The model suggests that intervening to improve these structures will generate greater benefit than addressing individual sub-cultures within

specific research teams, as important as these are.

The problem is configured as how to facilitate growth in a data driven economy. Core issues are reducing the costs of collaboration, clarifying property rights, developing efficient institutions, creating trustworthy environments, and being explicit about our value-base to the point of building this into data contracts. Figure 3 is a schematic of this model showing how a values-base and user-led innovation provide the context and engine underlying

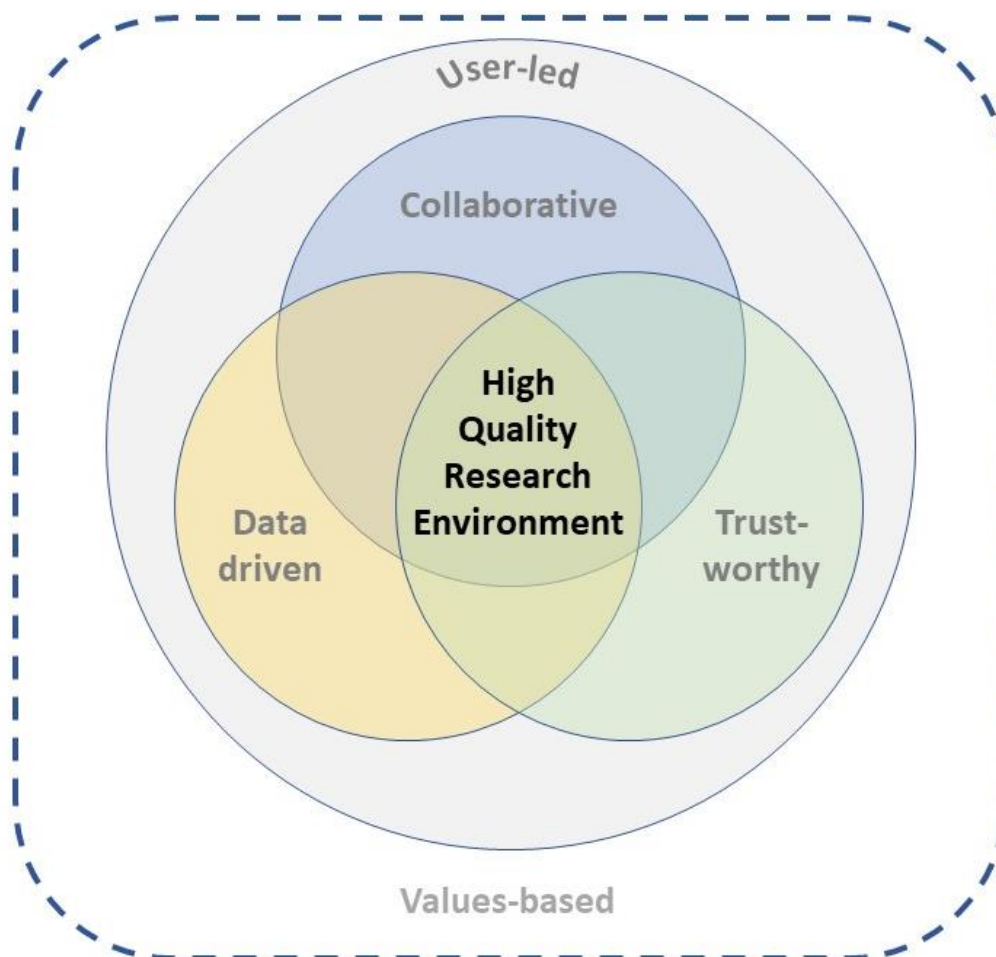


Figure 3:

A model for high quality research environments

collaboration, data access, and trustworthiness.

An implicit tension between competition and collaboration underlies the argument and shapes any scientific data economy. Competition between scientists and a judicious degree of private property rights provides a necessary incentive for the investment of personal energy, resources and funding into worthwhile projects, thus sustaining the quality of science. Competition between funders, in the sense of complimentary agendas, also helps maintain the quality of science, particularly in the originality of research. However, especially here, there is a substantial collaborative dividend available for synergising and not duplicating effort. Remove competition from either supply or demand sides, and the quality of funding calls and bids declines. However, to improve overall benefit from limited resources, aspects of competition need moderating.

To convey the argument, we have characterised this as a problem of club formation, since scientists by and large behave as though scientific data is a club good. One distinction of a good that is best organised in clubs is that the efficient quantity and quality of the good (research) is determined simultaneously with club size. This clearly happens routinely as scientists invite/join collaboration based on the likely outputs of a particular configuration. Our argument is that this type of market under-supplies science. In elaborating the argument, we point to platform technologies and collaborative incentives as market-corrective mechanisms likely to lower the transaction costs that hinder collaboration.

We argue that providing an efficient ecosystem in which these negotiations can take place requires the development of mature institutions around practice. These

institutions are not restrictive, but supportive of decentralised and spontaneous decision making, and provide the conceptual framework around which automated procedures that serve the interest of all stakeholders can be developed. Here, funders have a key leadership role to coordinate effort.

Finally, we focus on the value-base that underpins the science ecosystem; understanding the distinction between aesthetic and ethical judgement. We argue that promoting both motivates high quality science, by providing a framework for personal motivation, the efficient formation of clubs and the coalescence of clubs into 'super-clubs' for addressing specific scientific challenges. This framework rescues us from the trivial, helps us rise above self-interest, and provides the sense of common purpose that is characteristic of science at its best.

Funding:

Gallacher is supported by the UK Research and Innovation (UKRI) Medical Research Council (MRC) grant for Dementias Platform UK 2 [MR/T033371/1].

Acknowledgements:

The discussion that led to this paper started during research funded by UK Economic and Social Research Council's Transformative Research grant [ES/L003201/1] and currently continuing under HKSAR Research Grants Council, Collaborative Research Fund project number: C7105-21GF.

The authors would also like to thank Dr Sarah Bauermeister, Katherine Shepherd and Tim Munday for offering advice on the text and assisting with the preparation of the manuscript.

Ethics:

The work did not require separate ethics approval.

Competing interests:

Gallacher is Director of Dementias Platform UK, and was a member of the UK Biobank Steering Group between 2010 and 2022.

Webster declares no competing interests.

CRedit summary:

Gallacher and Webster contributed equally to conceptualisation, original draft writing, and review and editing of this manuscript.

References

1. REF 2021: [Results and submissions : REF 2021](#)
2. Munafo MS. Raising research quality will require collective action. *Nature* 2019;576:183.
3. [Research Culture - OD&PL Researcher Development and Research Culture \(leeds.ac.uk\)](#)
4. Research Culture: Embedding Inclusive excellence. Royal Society 2018
5. What researchers think about the culture they work in. Wellcome Trust 2020
6. Buchanan JM. An economic theory of clubs. *Economics* 1965;32(125):1-14.
7. Herbert et al. On the time spent preparing grant proposals: an observational study of Australian researchers. *BMJ Open*. 2013;3:e002800
8. Gabelica et al. Many researchers were not compliant with their published data sharing statement: mixed-methods study. *Journal of Clinical Epidemiology*. In press 2022
9. Carbon et al. An analysis and metric of reusable data licensing practices for biomedical resources. *PLoS One* 14 (3): e0213090.2018
10. Cornes, R. and Sandler, T. *The Theory of Externalities, Public Goods and Club Goods*. Cambridge University Press, New York. 1996
11. The Recovery Collaborative Group. Dexamethasone in Hospitalized Patients with Covid-19 — Preliminary Report. 2020 Jul 17 : *NEJMoa*2021436
12. Sherry et al. 1999. dbSNP: Database for single nucleotide polymorphisms and other classes of minor genetic variation. *Genome Research*. 1999;9:677-79
13. Bauermeister J et al. Landscape analysis for translational dementia data. In preparation
14. Bauermeister S. et al 2022: Research-ready data for multi-cohort analyses. *EJE* in press.
15. Gecko: [Genomics Cohorts Knowledge Ontology < Ontology Lookup Service < EMBL-EBI](#)
16. Dementias Platform UK Data Portal: <https://portal.dementiasplatform.uk/>
17. Bauermeister S. et al. 2020 The Dementias Platform UK (DPUK) Data Portal. *European Journal of Epidemiology* 2020;35:601-11. <https://doi.org/10.21203/rs.3.rs-937113/v1>
18. DP Australia: [Home | DPAU \(dementiasplatform.com.au\)](#)
19. ADDI workbench: [Alzheimer's Disease AD Workbench \(ADWB\) - ADDI \(alzheimersdata.org\)](#)
20. Wittgenstein L. Lecture on ethics. a1929