MOBILISING THE ACADEMIC STUDY OF ORGANISATIONS

Current insights and future possibilities for designing the institutions we need

Sir Geoff Mulgan

06 / 2026

WHAT IS TIAL

The Institutional Architecture Lab was formed in 2023 by Sir Geoff Mulgan, Jessica Seddon, and Juha Leppänen in an effort to help the institutional design community coalesce, learn together, and grow. Each of us has been involved in various stages of creating new organizations and other institutions. Like many other people, we have witnessed first-hand the absence of a formal community along the way — or a place where we can learn from past experience. We are aware that there is a lot of great work happening around the world, but nowhere to recognize it.

Authors

Sir Geoff Mulgan, *Professor of Collective Intelligence, Public Policy & Social Innovation at UCL; Co-founder at TIAL.*

TABLE OF CONTENTS

In brief

What insights do academic disciplines provide for the design of new organisations and institutions?[1] One of the goals of TIAL, the Institutional Architecture Lab, has been to link academic work on understanding organisations with the practical task of designing the next generation of public institutions, which we see as essential if the world is to thrive over the next few decades.

There is plenty of good work in many disciplines that provides insights into how organisations and institutions work, including research by several Nobel prize-winners. But there is a surprising disconnect between practice and research, little cross-pollination between the disciplines, and a strong bias towards analysis over design, and diagnosis over prescription. The position echoes the famous parable of the blind men and the elephant, with each discipline describing similar phenomena in radically different ways, and few offering more comprehensive or composite pictures.

The last few years have seen particularly impressive books analysing the importance of institutional innovation in the past from figures such as Harvard's Daren Acemoglu. But their writings have surprisingly little to say about the present or future. At a time of rapid change in organisational forms — much of it driven by technology — the academy has struggled to keep up both in terms of analysis and even more in terms of proposals and useful ideas. So, when decision-makers are tasked with creating a new institution, there is relatively little to draw from.

In this paper, I've tried to summarise the perspectives of different disciplines — economics, psychology, computer science, business

[1]There are many competing definitions that try to distinguish institutions from organisations. Given the nature of language this is not easy. We sometimes talk of the 'institution' of marriage, for example and some academic definitions go much further down this route, defining institutions to encompass all sorts of rules and norms (eg Douglass North, DiMaggio, Giddens and others do this), with the word organisations used to refer to entities like firms, public agencies or NGOs. Yet in everyday language, certainly in English, a very different approach is taken most of the time. Although people sometimes talk of things like marriage as institutions, the word is usually used to describe a subset of organisations — generally the ones that are more formal, official, and public in nature, including ministries, churches, public services, libraries, galleries, universities and so on. I have chosen to go along with the actual uses of the words in contemporary language. I'm sceptical of the value of using academic definitions which diverge too far from the vernacular.

studies, organisation studies, political science, history, law, international relations, anthropology, design, and complexity. In each case, I make short suggestions on what would be useful from each, before turning to what a more synthetic approach might look like, in particular, using insights from biology and computation to see organisations as living things and addressing the dynamics of ecosystems of organisations that compete and cooperate.

The paper asks of the people working in academic disciplines: how are you engaging with, and learning from, other disciplines? And how could your knowledge be useful to a world that badly needs to reform its public institutions?

1. Insights from the major disciplines

This section sketches how existing disciplines look at organisations and institutions of all kinds — what they see and highlight, and how they think.

Economics: incentives, markets & efficiency

Economics offers some of the most systematic frameworks for thinking about institutional design. Its central contribution is the analysis of incentive structures: institutions reliably produce the behaviours they reward, not the behaviours they intend. This distinction helps to explain many institutional failures. Economists warn of principal-agent problems — when those who run institutions (agents) have interests that diverge from those they serve (principals) — and of information asymmetries, where different parties know different things, enabling gaming and capture.

Public choice theory adds a sobering corrective to idealism: public officials and bureaucrats are self-interested actors just like anyone else, not benevolent social optimisers. Institutions must therefore be designed as if populated by ordinary, self-seeking people. Regulatory capture — where regulated industries come to dominate their regulators — is a perennial hazard to guard against.

The economics of collective action highlights why voluntary cooperation fails in the provision of public goods: free-rider problems mean that without compulsion or clever incentive design, public institutions will be underfunded and underused. Elinor Ostrom's Nobel-winning work on 'governing the commons'

demonstrates, however, that communities can self-organise effective institutions without either markets or states when rules are locally designed and enforced. Cost-benefit analysis, market-design techniques, and the insights of the growing field of mechanism design are practical tools that economists contribute to institutional architects.

What might we want more of from economists? We could hope for better understanding the full range of economic organisations, from mutuals and coops to DAOs; better understanding ecosystems of organisations, and coupled systems (e.g. where states intervene in crises, for example, to enable business to thrive); better understanding the economics of intelligence design in organisations, and trade-offs of investment in observation, memory, creativity etc; and better use of economic history and economic sociology, as there is relatively little cross-pollination between the different strands of economic thinking.

Psychology: behaviour, bias & motivation

Psychology shows that the humans who populate and use institutions are not the rational maximisers that economics often assumes. Cognitive biases — systematic errors in judgment documented by behavioural economists such as Kahneman and Thaler — profoundly affect how institutions are used. Anchoring, availability bias, status quo bias, and loss aversion all shape how citizens and officials respond to institutional rules and defaults, though others have convincingly shown the blind spots and errors of this approach, emphasising instead the effects of systems and why some of these apparent biases are in fact quite rational (for example, in the work of Gerd Gigerenzer and Piero Luigi Sacco).

Nudge theory, drawing on behavioural insights, argues that the architecture of choices within an institution matters enormously. Defaults, framing, salience, and social norms can steer behaviour far more cheaply and effectively than coercion or financial incentives. The UK's Behavioural Insights Team ("Nudge Unit") has demonstrated this across tax compliance, pension enrolment, and public health, though there are many sceptics and critics (some

captured in the recent book by Nicholas Chater and George Loewenstein).

Motivational psychology distinguishes intrinsic from extrinsic motivation, warning that over-reliance on financial incentives can "crowd out" the professional pride, civic duty, and intrinsic satisfaction that actually drive much public-sector performance.

Organisational psychology highlights the importance of psychological safety — the belief that one can raise concerns without punishment — for institutional learning and error correction. Trauma-informed design is emerging as a framework for institutions serving vulnerable populations, ensuring that procedural experiences do not re-traumatise those who need help most.

What might we want more of from psychology? One priority is more non-WEIRD psychology (to use Joseph Heinrich's formula), for example, understanding family businesses or different state forms around the world. We could benefit from a better grasp of cognitive processes in organisations, including error, collective mental illness, illusions, and more insight into the psychological dynamics of organisational failure and collapse — what happens when people lose faith.

Computer science: systems, data & digital infrastructure

Computer science brings a systems-thinking perspective to institutional design, treating institutions as information-processing architectures. Concepts from software engineering — modularity, abstraction, versioning, and testing — translate surprisingly well to institutional design. Modular institutions with clearly defined interfaces between components are easier to reform, audit, and upgrade than monolithic ones.

Computer science focuses on decisions and information — inputs, processing, memory, decisions, and learning, and concepts such as updating (from the overlap of neuroscience and computer science).

The idea of stacks, both in their narrow interpretation (the OSI model, Internet, and examples like the India Stack) and in a broader sense of using communications to enable more layered design of institutions, are all relevant to the late 2020s.

Digital public infrastructure is now one interesting foundational idea, pointing to the potential for protocols and stacks to guide design. Another key frame is data, with institutions designed around data governance from the outset: who collects what, on what legal basis, with what retention periods, subject to which rights of access and correction. The GDPR framework and the open data movement represent competing institutional logics — protection versus transparency — that must be consciously balanced.

Algorithmic decision-making raises new questions. Some concern standards and how these are set. Another concerns institutional accountability. When an institution delegates decisions to an algorithm — in welfare benefit allocation, sentencing recommendations, or credit scoring — traditional accountability mechanisms (appeal, judicial review, parliamentary scrutiny) struggle to function. Explainability, auditability, and contestability must be built in from the start. Cybersecurity is an institutional design challenge, not merely a technical one: institutions are attacked through their people and processes as much as their systems. Platform governance — how digital platforms moderate speech, allocate attention, and set rules — has itself become a major domain of public institutional design.

What might we want from computer science? One obvious priority is better theories and practice of how the human and computational combine — computer science has been surprisingly slow to absorb broader insights about how intelligence functions in the real world. Linked to that is better orchestration of evidence on AI in organisations, and understanding organisations as assemblies of human and machine intelligence. This will become increasingly vital for the work of design.

Business studies: strategy, management & performance

Business studies — encompassing strategy, management, and operations — contribute practical tools for institutional design, though they must be applied with care in public contexts. There are many frameworks from business, ranging from platforms to franchising and licensing, multi-divisional and matrix structures.

Strategic management frameworks like SWOT analysis, stakeholder mapping, and theory of change help institutions clarify their purpose and map the environment they operate in. The discipline of defining a clear mission and refusing mission drift is one of business's most transferable lessons.

Performance management offers both promise and peril. Key performance indicators (KPIs) can focus attention and enable accountability, but "teaching to the test" — Goodhart's Law — means that measures become targets and cease to be good measures. The NHS and education systems have provided cautionary evidence of how league tables and targets distort behaviour. The balanced scorecard approach attempts to address this by measuring multiple dimensions simultaneously.

Supply chain thinking and operations management apply to public service delivery: institutions must design workflows, manage queues, reduce bottlenecks, and build resilience into their delivery systems. Lean management, adapted from manufacturing, has been applied to public administration to reduce waste and improve flow.

Leadership studies highlight the importance of institutional culture — "the way we do things around here" — as a determinant of performance that resists top-down reform, often providing detailed insights (e.g., in the work of Mintzberg). Change management frameworks (Kotter, ADKAR) offer structured approaches to institutional reform.

What might we want from business studies and schools? We could certainly benefit from a better understanding of emerging practices, e.g., the ultra-flat models of firms like Nvidia (with large

numbers of direct reports and very different leadership styles); platforms; the rapidly evolving governance structures adopted by very big firms like OpenAI and Anthropic; new patterns of licensing and franchising; and the different approaches of the social economy. Tools such as the business model canvas could be usefully adapted to institutional design.

Organisational studies: structure, culture & learning

Organisational studies is perhaps the most directly applicable discipline to institutional design. It examines how formal structures (hierarchies, matrices, networks), informal cultures, and professional norms interact to shape what institutions actually do — as opposed to what their charters say they should do.

A central insight is the distinction between formal and informal organisation. The real institution is always richer, stranger, and more interesting than the org chart. Informal networks, professional identities, subcultures, and tacit knowledge determine how rules are interpreted and whether change takes hold. Reformers who attend only to formal structure routinely fail because they leave informal systems untouched.

Institutional theory (DiMaggio and Powell) explains why organisations in the same field tend to become structurally similar over time — a process called isomorphism — through coercive pressures (regulation), mimetic pressures (copying successful peers), and normative pressures (professional standards). This helps explain why new institutions tend to resemble old ones even when designed to be different. Learning organisation theory (Senge, Argyris) distinguishes single-loop learning (fixing errors) from double-loop learning (questioning the assumptions that produced the errors) — the latter being essential for institutions operating in complex, changing environments. High-reliability organisation research examines how nuclear plants, air traffic control, and aircraft carriers achieve very low error rates in high-stakes environments, with lessons for institutional design around safety and resilience.

Culture theory and grid/group analysis (drawing on the work of Mary Douglas and Michael Thompson) suggests that all real organisations are compromises between hierarchy, individualism, egalitarianism, and fatalism, and all tend to fail if any becomes too dominant.

What might we want from organisational studies? More interest in design as opposed to analysis; better quantitative methods to understand ecosystems of organisations, patterns of birth and death, and evolution; and more engagement with technology, which is still often missing from much academic work.

Political science: power, legitimacy & governance

Political science is concerned with the most fundamental questions of institutional design: who has authority, from where does it derive legitimacy, how is power checked, and who is included in or excluded from governance? These questions precede all technical choices about structure and incentives.

Political science tends to focus on the specific conditions in which new institutions are formed and how these reflect the power of different interests. Constitutional design — the architecture of separation of powers, federalism, and rights — establishes the environment in which public institutions operate. Institutions that are not embedded in a coherent constitutional framework are vulnerable to capture, abuse, and collapse. Checks and balances are not bureaucratic friction to be minimised; they are the technology by which power is kept accountable. There is also a strand of work on pluralism: how democracy requires counters to what Robert Dahl called 'closed hegemony' that drowns out competing voices.

Legitimacy theory distinguishes input legitimacy (is the institution democratically authorised?), throughput legitimacy (are processes transparent and participatory?), and output legitimacy (does it deliver results?). New public institutions require attention to all three dimensions. Electoral politics shapes institutional incentives profoundly: short electoral cycles, partisan pressures, and

ministerial churn all undermine institutions' capacity to pursue long-term goals. Independent institutions — central banks, statistical agencies, regulatory bodies — represent an institutional design response to this problem, trading democratic immediacy for credibility and continuity.

Deliberative democracy theory argues for designing institutions with genuine public participation and structured deliberation, not mere consultation. Citizens' assemblies, as used in Ireland and the UK, are an emerging institutional form that has produced politically durable outcomes on otherwise intractable questions.

There is also a long strand of work in political science which emphasises difficulty: wicked problems, irrational processes, impossible jobs and intractability, with some claiming that there are long-run trends towards ungovernability in conditions of permacrisis. These theories often struggle to explain the long history of policy successes and can be ahistorical, ignoring past periods that saw multiple crises simultaneously, and they are not so useful for the practical work of institutional design.

What might we want from political science? More curiosity to learn from other disciplines and interest in the options for design, as well as more precision on the role and limits of path dependence; more engagement with the ways that technologies shape the options for policy and administrative design.

History: evolution, path dependence & failure

History is the laboratory of institutional design, providing the empirical record from which all other disciplines draw their examples. Its central methodological contribution is a warning against overconfidence: almost all new institutions fail in ways their designers did not anticipate, and history is littered with earnest institutional experiments that produced catastrophic unintended consequences.

Path dependence is one of history's most important conceptual contributions. Institutions acquire constituencies, vested interests, and embedded practices that make them extraordinarily difficult to

change even when their original rationale has evaporated. The QWERTY keyboard and the gauge of British railway tracks are classic examples; in public life, the survival of institutions long past their useful life — from parliamentary procedure to welfare eligibility rules — follows the same logic. Design choices made early become almost impossible to reverse. Conversely, many long-standing institutions become interwoven with everyday life and mentalities.

The comparative historical record reveals recurring patterns of institutional success and failure. Successful institutions tend to have clear mandates, insulation from short-term political pressure, adequate resourcing, and mechanisms for self-correction. Failed institutions typically suffer from mandate overstretch, underfunding, capture by narrow interests, or loss of public legitimacy. The history of welfare state construction, colonial administration, post-conflict state-building, and development banking all offer rich and often cautionary material.

A distinctive contribution has come from economic historians who tend to take a very different view from economists, often putting much more emphasis on the roles institutions play. The Nobel Prize for economics in 2024 was awarded to Daren Acemoglu and colleagues for their detailed account of the role of institutions in shaping national prosperity and routes to societal success. These emphasise the unlikelihood of these paths (though, as indicated earlier, they offer few clues to current design).

Finally, history reminds institutional designers that the "problems" they are addressing are rarely new. Someone has usually tried before. Understanding why previous attempts failed is the most reliable preparation for doing better.

What might we want more of from history? Better integration of quant and qual methods to study long-term patterns, including understanding of punctuated equilibrium, using cliometrics, etc., and interest in the reappearance or revival of organisational forms.

Law: rules, rights & accountability

Law and legal studies provide the constitutive framework within which all public institutions exist: without legal personality, mandate, and powers conferred by statute or constitution, there is no institution at all. The first contribution of legal thinking is therefore foundational — getting the enabling legislation right matters enormously. Vague statutory mandates invite mission creep and judicial challenge; overly narrow ones prevent adaptation to changing circumstances.

Administrative law is the branch most directly concerned with institutional design. Its core preoccupations — procedural fairness, the duty to give reasons, proportionality, judicial review, and the rule against bias — are not merely legal technicalities but embody deep principles about how public power should be exercised. Institutions designed without administrative law expertise routinely create decision-making processes that are procedurally defective, generating litigation and eroding legitimacy.

Rights-based frameworks — drawn from constitutional law, human rights instruments, and equality legislation — impose substantive constraints on what institutions may do and affirmative obligations about how they must treat those they serve. Designing institutions to be rights-compliant from the outset is far cheaper than retrofitting them after adverse judgments.

Regulatory law contributes to the concept of enforcement pyramids: effective compliance is achieved not through maximum punishment but through graduated responses that deploy persuasion, warnings, civil penalties, and criminal prosecution in ascending order of severity. Legal pluralism — the recognition that multiple overlapping legal orders (state law, customary law, religious law, professional norms) simultaneously govern behaviour — is critical when designing institutions for diverse or international contexts. The enforcement pyramid concept, in particular, is a practically important idea that pure economics or

political science tends to miss. A legal perspective helps to explain the striking isomorphism of organisational forms.

What might we want more of: lessons from new legal forms such as CICs, BCorps, DAOs, and Public Benefit Corporations? And more imagination in how law could be evolved, for example, for transnational organisations, or the use of blockchain and other technologies.

International relations: norms, regimes & global governance

International relations, or IR, offers insights that matter even for purely domestic institutions, because no institution operates in isolation from global pressures, norms, and interdependencies. It adds the ecosystem perspective: how does this new institution fit within the wider architecture of existing bodies, and how will powerful actors exploit any gaps or inconsistencies between them? It also introduces the constructivist insight that institutions don't just change behaviour — over time, they change identity and preferences, which is a much deeper and more durable form of influence.

Its most direct contribution is the study of international institutions themselves — the UN, WTO, IMF, WHO, and hundreds of other bodies — as a rich empirical laboratory of what makes supranational governance work or fail.

Regime theory examines how international regimes (sets of principles, norms, rules, and decision-making procedures) emerge, persist, and decay. The conditions for regime formation — concentrated benefits, clear rules, credible enforcement, and shared understanding of the problem — translate directly to domestic institutional design. Liberal institutionalism argues that institutions reduce transaction costs, provide information, and create focal points for cooperation that would otherwise be impossible; these mechanisms operate at domestic as well as international scales.

Constructivism contributes a focus on norms and ideas: institutions do not merely constrain behaviour, they constitute actors' identities

and interests over time. International human rights norms, for instance, have gradually reshaped what states believe they are entitled to do to their own citizens. This socialisation function of institutions — their capacity to change, not just channel, the preferences of those within them — is easily underestimated in purely rationalist frameworks.

IR also contributes to the analysis of institutional fragmentation and forum-shopping: when multiple overlapping institutions govern the same domain, powerful actors exploit the gaps and inconsistencies between them. Coherent institutional architecture — at domestic and international levels — requires deliberate attention to how new institutions fit within the existing ecosystem.

What might we want? IR has largely lost interest in options for the future and the design of new institutions for space, oceans, AI, and warfare, or the organisation of global intelligence.

Anthropology: culture, meaning & local knowledge

Anthropology is the discipline most attentive to what institutions actually mean to the people who encounter them, and to the deep cultural assumptions that institutional designers unwittingly embed in their creations. Its primary methodological contribution is ethnography — sustained, immersive observation of how institutions function from the inside — which consistently reveals gaps between official accounts and lived reality that other methods miss. It asks: "But what does this look like from the inside, at ground level, to the people actually living with it?

The anthropology of bureaucracy (drawing on scholars like David Graeber and Akhil Gupta) documents how formal rules are translated, distorted, resisted, and repurposed by street-level bureaucrats and citizens. The discretionary power of frontline officials — those who actually deliver institutional outputs — is far greater than institutional designers typically assume, and the exercise of that discretion is shaped by local culture, professional norms, and informal hierarchies.

Cultural translation is a persistent challenge. Institutions transplanted from one cultural context to another without adaptation routinely fail. The history of development institutions imposing Western governance models on societies with different kinship systems, property concepts, and conceptions of authority is a long catalogue of instructive failure. Anthropology insists that local knowledge — what James Scott calls mētis, the practical wisdom embedded in communities — must be incorporated into institutional design, not displaced by it. He showed in his classic Seeing Like a State how institutions represent and simplify social reality, always at a cost, and made the case for resisting generalisation and standardisation.

Ritual and symbolism matter more than rationalist frameworks acknowledge. Institutions communicate meaning through their architecture, ceremonies, uniforms, forms, and spatial arrangements. These symbolic dimensions shape how institutions are perceived, trusted, and used — and their neglect produces institutions that are technically correct but culturally alien to those they are meant to serve.

What we might want more of from anthropology: reflection on the limits of lived experience, or how structures shape mentalities, and more collaborative work with disciplines that take a very different view.

Design: co-production, affordances, systems, and prototypes

The study of design has provided many insights into how institutions work and could work differently. Herbert Simon's "Sciences of the Artificial" distinguished natural sciences (describing what is) from design sciences (prescribing what ought to be) and argued that administration and policy are fundamentally design activities. He hoped for a rational and scientific approach to every aspect of governance, including computation and modelling. However, these hopes were not realised. A few years later, the work on wicked problems (Rittel

and Weber) problematised this approach, arguing that social and policy problems resist the linear, engineering model of problem-solving — they are "wicked" because defining the problem is the problem. Public institutions are constitutively engaged with wicked problems (poverty, crime, health), and this insight argues against technocratic, blueprint-style institutional design in favour of iterative, adaptive approaches.

A later shift moved from design for users to design with users and also had clear institutional implications. Scandinavian participatory design traditions (emerging from labour movements in the 1970s) argued that those affected by a system should shape it, echoing ideas from philosophy and political science on deliberative democracy (Habermas, Fishkin, Landemore) and seeing institutions as forums for structured co-creation of public decisions, mirroring the interest in co-production in public services (Ostrom, Bovaird) with citizens as active participants in service design, not passive recipients. Design absorbed the critique that top-down institutional design systematically excludes the tacit knowledge of those closest to a problem.

Design also provided useful insights into the role of symbolic design — the architecture of courts, the rituals of parliament, the aesthetics of public spaces, which all communicate authority, accessibility, neutrality, and power: indeed, no institution can be understood separate from the fictions or claims that it makes. James Gibson's concept of affordances (what an environment makes possible or invites) can be applied to institutions: institutional rules and structures afford certain behaviours and constrain others — not through command but through the shape of the environment. Behavioural design (Kahneman, Thaler & Sunstein's various approaches to choice architecture) is an explicit application: designing choice environments to channel behaviour without coercion, connecting to older work in political science on how constitutional design shapes political behaviour (Lijphart, Tsebelis).

Design culture's embrace of prototyping, drawn from civil society and software — building cheap, testable versions before committing — is a corrective to the grandiose, irreversible institutional reforms that characterize much political change, but is challenging in law-based contexts. It overlaps with experimental

methods such as the use of randomized controlled trials and policy labs (Nesta, MindLab, various government innovation units), which have tried to bring this ethos to public sector reform. These suggest the concept of minimum viable institutions — starting small and scaling what works. Design's normalization of failure as informative rather than shameful is a genuine cultural challenge for public institutions with accountability pressures.

Design has increasingly absorbed ideas from systems thinking (Beer, Meadows, Forrester), which is directly applicable to institutions: seeing how institutions generate feedback loops — perverse incentives, goal displacement, Goodhart's Law; the role of anticipation in uncertainty; attention to second and third-order effects rather than optimizing local components; and the concept of leverage points, places in a system where a small intervention produces large change — is highly useful for institutional reform strategy.

Design has had a significant influence on public administration over the last few decades. However, this has primarily focused on services, promoting distinct mindsets and ways of thinking. Whereas service design and digital design have developed specific methods, there isn't yet anything quite comparable for institutional design, and Herbert Simon's hopes sixty years ago for a comprehensive approach linking design to governance have yet to be realised.

Complexity: cybernetics, systems, and the dilemmas of emergence

The many currents of research on systems, cybernetics, and complexity have had much to say about how organisations work, from Stafford Beer, Donatella Meadows, and Peter Checkland to Stuart Kauffman. Much of this work emphasises limits: what can't be known or planned; the virtue of improvisation (making paths by walking them), which implies not working too hard to design or specify the architecture of organisations; attending to the broader systems within which organisations operate; and addressing emergence, non-linearity and contradiction as inevitable in human

systems, all of which call into question the more rationalist models of organisational design and analysis.

One stream of work, that of cybernetics from the 1940s onwards, emphasised how organisations regulate themselves through negative feedback (corrective) and positive feedback (amplifying). This reframed management as a process of sensing and responding to deviation, not just commanding. Ross Ashby's law of requisite variety became popular as a frame, arguing that a controller must have at least as much variety (or range of responses) as the system it governs: you cannot manage complexity with simplistic rules. This theory has struggled in practice, since it is perhaps never feasible (and all organisations have little choice but to use methods far simpler than the environments they are trying to influence).

Stafford Beer's Viable System Model applied cybernetics to organisations directly, showing that viable organisations need five nested subsystems (operations, coordination, control, intelligence, and identity/policy), with each subsystem fractal (and potentially containing all the others). Like Ross Ashby, he emphasised amplifying internal diversity to enable learning in ways that match the complexity of the external environment. His ideas also struggled when put into practice.

Drawing on biology (von Bertalanffy) and later sociology (Parsons, Luhmann), systems theories were another challenge to overly reductionist, mechanistic views of organisations. These tend to emphasise the ways that organisational properties (culture, strategy, innovation) emerge from local interactions and cannot be straightforwardly designed top-down; how relationships between components matter as much as the components themselves; and how organisations can be understood as open systems, open to their environments, exchanging energy, information, and resources.

A parallel strand of work looked at organisations as sociotechnical systems (for example, the work of Trist & Bamforth at the Tavistock Institute), which means that both their social and technical subsystems need care and attention. Tavistock researchers also looked at organisations not as logical entities pursuing defined goals but as emotional systems driven by unconscious forces. At the heart of their practice is systems

psychodynamics, the study of how unconscious processes shape organisational life, with social structures often serving as defences against individually experienced anxiety, guilt, and doubt. Wilfred Bion argued that groups simultaneously operate on the conscious level (the "work group"), which is focused on the explicit task or goal, and the unconscious level (the "basic assumption group"), which is driven by unspoken assumptions and emotions. He identified what he called a "hatred of learning by experience", as groups unconsciously cooperate to avoid dealing with their real tasks and resist the uncertainty involved in learning about themselves. Like some of the other theoretical approaches described above, they argued that the best way to understand a system, or an organisation, is to try to change it.

Senge used several of these traditions to feed into his ideas about "learning organisations" that can avoid being trapped in toxic or unproductive patterns while the theories of Complex Adaptive Systems (CAS) similarly emphasise self-organising systems of agents following local rules, with order arising without central coordination, or even that organisations thrive best on the edge of chaos, on the cusp between rigid order and complete disorder. Too much control kills adaptability; too little produces incoherence.

Other relevant tools include the analysis of non-linearity and sensitivity, which show how small interventions can have disproportionate effects (and vice versa), which undermines linear planning models and highlights the limits of prediction. The development of models analysing fitness landscapes provided another way to see organisations as continuously adapting to shifting competitive environments. What was an optimal strategy can become a trap.

These many theories have all struggled to generate either falsifiable hypotheses or predictions, or much that can be quantified, and some have been hard to operationalise without the help of very skilled facilitators. Their other challenge is that the emphasis on emergence, and the impossibility of design, clashes with the fact that most organisations are in fact designed — with legal forms, organograms, and proposals used to generate funding or investment. Their reality may diverge from the plans. But they don’t simply emerge.

However, these theoretical traditions are richly suggestive and influenced the development of agile and lean methodologies, with iterative, feedback-driven, decentralised, network and flat organisational structures. They also encourage organisations to reduce hierarchy so as to increase variety and responsiveness and improve organisational learning, with reflexive capacity rather than fixed procedures. The key legacy of these traditions is a way of thinking, seeing organisations not as machines to be engineered but as living, adaptive systems to be cultivated.

2. Where next?

If section 1 mapped out the current state of academic research, where might we hope for research to head in the future? How could more synthetic work be done, integrating insights from multiple disciplines? And are there new intellectual frameworks that might help to give some shared direction to these?

Syntheses and overlaps

There are many overlapping ideas and focuses in these disciplines. Economics often collaborates with law, emphasising formal structures and incentives; behavioural design in psychology feeds into organisational studies, often emphasising that formal rules and structures may matter less than the lived experience of the people inside and around an institution, meaning that it's vital to get the informal right as well as the formal.

Some emphasise individuals (including most North American psychology and economics) while others (in particular social psychology and anthropology) emphasise a group mind, identity, solidarity, and shared cultures. Political science and its traditions of institutionalism overlap with history, emphasising path dependence and the specific conditions in which new institutions arise (though all recognise that it's risky to identify too clearly any 'lessons from history').

Some disciplines highlight key tensions and ambiguities. Political science highlights a fundamental tension between democratic accountability (short cycles, political responsiveness) and institutional effectiveness (long time horizons, insulation from pressure). Every public institution must consciously navigate this trade-off. Computer science is increasingly relevant here too, as algorithmic delegation of decisions creates accountability gaps that no existing institutional form fully resolves, but despite several decades of work on human computing interfaces, and the various contributions of 'Actor Network Theory', we still lack disciplinary frameworks that are comfortable seeing organisations

as combinations of human and machine. The systems approaches aspired to synthesise others and can provide a common language from which to draw ideas from other disciplines.

From another perspective, Elinor Ostrom and colleagues attempted a synthetic approach in the 1980s, partly drawing on game theory, which led to the 'Institutional Analysis and Development' (IAD) models that tried to link institutional design to the nature of outputs, contexts, information, incentives, and authority. These aimed to understand the different institutional designs needed at operational levels, collective-choice or policy levels, and at a constitutional level. At one point, it was hoped these could guide the design of new institutions to manage topics such as data. However, these methods have not turned out to have predictive power (as was originally hoped), to be as widely applicable, or to be amenable to coding, and they can be oddly static.

Future options

TIAL's work so far has drawn on methods from design, innovation, and social movements to suggest alternative options for organisational and institutional design. It emphasises the vast range of possibility spaces for the future and the need to resist organisational isomorphism — from models like BCorps and CICs to DAOs, complex webs of licensing to platforms — and options inspired by nature, such as myceliums or octopuses, and from computer science (meshes and stacks).

This work has shown that some key insights may come from disciplines that have had little impact on organisational thinking in the past, in particular biology and the study of collective intelligence. Here I highlight five potential future directions, hoping to prompt others to prose their alternative proposals.

Organisations as forms of artificial life

There is a good case for seeing institutions as living things, and as examples of 'symbio-genesis': the evolutionary mechanism where

new species, organs, or complex traits arise from the long-term, permanent integration of different organisms.

Organisations are often assemblies of this kind, and they have patterns of birth, reproduction, life, and death which mirror living organisms. They have metabolisms. And they have boundaries and membranes. They are also artificial life forms, not unlike material technologies: they are designed, created, and are at least partially programmable.

- Like living things, they tend to have a commitment to their own survival as paramount, so that ends become means (autopoiesis), overriding the predictions of other disciplines (economics, psychology, etc.).
- They operate within ecosystems of competition and cooperation that profoundly shape what they are and how they act.
- As with living things, organisational forms that can reproduce are, paradoxically, more dynamically stable than non-living things.
- Within each organisation, like cells, the whole history of organisations is visible. Just as human biological history is visible in the cell, from the salty water within it to the mitochondria, so many organisations contain traces of past organisational forms, family, tribes, monarchy, feudal, capitalist, egalitarian....

The idea of institutions as organisms is very old — Thomas Hobbes, for example, wrote of the state as 'an artificial man' fuelled by sovereignty, which was an 'artificial soul'. The firm is an artificial person, able to make commitments that far outlast the individuals who make it up. The trust is an artificial organisation imprinted with values and purposes. Indeed, one measure of progress in a complex society is the range and diversity of these artificial forms, with their impersonal personalities, and we should expect a proliferation of forms as agentic organisations involving ever less human agency spread.

Organisations as systems of thought

A second related approach emphasises that the organisation of intelligence is a key aspect of life, both in living organisms and in organisations: observing, remembering, creating, with cognition at the core of all institutions, and needing substantial inputs of scarce energy. Indeed, it's fair to say that any organisation is a kind of artificial intelligence: it is certainly artificial and made, and it only exists as a way to orchestrate intelligence beyond the limitations of individual brains.

This lens may be equally fruitful and complementary:

- Analysing organisations as cognitive systems, with heuristics, algorithms, models, dedication of resources to tasks such as memory, creativity, prediction, and the use of 'triggered hierarchy' to handle decisions.
- Using variants of assembly theory which emphasise these combinations and path dependence, exploring possibility spaces.
- Looking at how all organisations generate representations of the world and fictions with dynamics that tend to take these further from reality until crises bring them back.
- And studying how organisations invest in metacognition to avoid the pitfalls of process, groupthink, myopia, and more.

Ecosystems of organisations

Perhaps the most obvious gap in all of the disciplines sketched earlier is the lack of a theory to understand how ecosystems of organisations work. All organisations exist in cooperation and competition with other organisations. But this is poorly understood. Business and economics have some well-developed theories of competition and market dynamics, and sub-disciplines concerned with the organisation of supply chains, and much

progress has been made using network analytic methods to understand the relationships between businesses — collaborations on research, products, manufacturing, services, marketing, and more. But we still lack a good theory of the interaction between businesses and states (which play decisive roles in economic crises) or with regulators (other than fairly simple theories of capture).

There is little or no good macro theory of how the vast number of organisations in a society interact — from small firms to co-ops, primary schools to charities, associations to multinationals. Such study would look in more depth at the dynamics of creation and destruction; how new forms mutate and spread; the role of recessions, droughts and famines in culling some kinds of organisation, while others prove more resilient; how every organisation seeks to improve its operating environment through advocacy and lobbying; whether there are advantages in pluralism of organisational forms; the velocity of systems (speed of emergences and disappearance); the power laws of organisational size; the shape of different sectors, and the mix of big, medium and small; how slow organisations, grounded in land, money or traditional status interact with very fast ones.

There are many potential analytic tools to be adapted from the study of natural ecosystems, which could be fruitful for understanding these dynamics and their influence on the real life of organisations, and the many tools of computational social science make these issues more tractable than in the past.

Fast and slow: aligning organisations and environments

A key issue for organisational design is how to handle speed. Some organisations move slowly and live for a very long time — the Catholic Church, some universities, and foundations are all examples. Others are more like butterflies, thriving briefly but brightly.

In principle, any ecosystem needs a mix of different organisational timescales for purposes as varied as religion, learning,

infrastructures, pensions, agriculture, forestry, and software, just as a forest contains organisms with widely varied timescales of life and action, from tall trees living for centuries to small insects lasting for days. But we lack any good theory to make sense of what a healthy balance is.

Meanwhile the obvious challenge of the contemporary world is an operating environment that moves far faster than the institutions meant to govern it, with many governments struggling to work out how to use or regulate technologies as they also face a rise in the number of emergencies they have to handle: climate, heat, floods, cyber, terror, which share some features but are also different (so that the next pandemic is likely to be different from the last one).

The mismatch between capabilities and tasks is glaringly obvious in regulation and law in relation to technology, but also affects culture, daily life, trade, and more. The mismatches become visible in many forms: the half-life of apparent successes may shrink; contradictions multiply (like officials asked to be rigorous in managing risk but also creative and entrepreneurial); and, because of asymmetries of knowledge, the gains to businesses that can achieve goals through lobbying as opposed to better products and services increases (hence the explosion of numbers of lobbyists in Washington and Brussels).

These patterns raise the question of what insights from academic disciplines can provide guidance as to how, through institutional design, to balance speed and slowness, and appropriate response to novelty without hysterical over-reaction. This matters less for big companies, which can afford high levels of risk and volatility, than it does for public institutions, which sometimes have a role of dampening change, taking the long view, ensuring that decisions are not made in an eternal present, but in other cases — such as war — have to move very fast. In the case of fast-changing technologies, there is a clear need for institutions with powers and capabilities to act very fast, and others with a design for much longer time horizons.

Organisations produce mentalities that produce organisations

The co-evolution of ways of thinking and organisations is a fascinating aspect of everyday life. Hierarchies need hierarchists, and markets need individualists: they attract people with dispositions but then shape and sharpen those dispositions too. The same is true of roles within organisations — safety managers, auditors, product developers, designers: each requires a certain mindset as well as technical skills. Some theorists in the past (such as Jane Jacobs and Mary Douglas) emphasised the need for complementary cultures and mindsets (and the pathologies that arise when these overreach), and every real organisation needs very different mentalities for distinct roles. Yet the mapping of these remains relatively underdeveloped.

3. Theories for design and action

The emerging lexicon of institutional design options that TIAL has developed — with concepts such as the mesh, stacks, Interchanges, mycelium, or octopoid organisations — provides a good challenge for each discipline to provide insights into their potential virtues and vices.

But this work also highlights a gap in the work of academic disciplines that remain predominantly inward-looking, with not much systemic analysis, few theories with any predictive power, or much healthy feedback between diagnosis and prescription.

Sharpening up the combinations of theory and practice is becoming a priority for the next few years, and confidence in public institutions is decaying. But it remains unclear which universities, philanthropic foundations, or other centres of research currently have either the capacity or interest to act to fill the gaps.

THE INSTITUTIONAL ARCHITECTURE LAB

Get in touch with us for advisory or speculative work
around institutional architecture

TIAL.ORG