Running head: COLLECTIVE INTELLIGENCE

Collective Intelligence in *Medias Res*:
Mulgan, Warfield, and the Collective Intelligence Network Support Unit (CINSU)

Michael J. Hogan

Address correspondence to:

Michael J. Hogan, PhD,

Department of Psychology,

NUI, Galway,

Ireland.

Email: michael.hogan@nuigalway.ie

Collective Intelligence in *Medias Res*: Mulgan, Warfield, and the Collective Intelligence Network Support Unit (CINSU)

There are many ways of thinking about human intelligence. The humorous quip you might hear on the street is that academics have generated so many ways of thinking about intelligence that the concept is now completely unintelligible. Still, talk of intelligence has not yet gone out of fashion. I recall sitting opposite a famous Scottish psychologist in an Edinburgh café, asking the question, *How would you define intelligence?* He answered succinctly with a smile: *Intelligence is the ability to figure things out.* His definition was as much an invitation to explore as anything else.

It's in our nature to explore. Throughout human biological and cultural evolution – and the slow process of learning how to speak, write, draw, count, compute, and fashion tools to support individual and collective action – intelligence has always involved an ability to *figure things out* and adapt to challenging environments. Migrating out of Africa 50,000 years ago to inhabit every corner of the globe today was no simple task for *Homo sapiens* – they had to figure out lots of things along the way. Similarly, global life expectancy has more than doubled from 1900 to 2018, and to achieve this *Homo sapiens* had to figure out a lot of things about how to stay alive for longer.

Intelligence, when it manifests, converges on some decision or action that is open to observation, reflection, and, sometimes, revision and experimentation. When we figure something out, we usually have something to show for it. An idea is voiced, a tool is modified, a group sets off in a particular direction across the savannah. Of course, decisions and actions accumulate over historical time. The more we have *figured things out* over time the more changes we have made to our environment. Our environment has become more complex, in the sense that it now includes many artefacts of human decisions and actions -- all the *products* of our past intelligence -- from languages to lampposts, empires to elementary schools, bridges to computers, saucepans to social media.

Homo sapiens are increasingly aware of their own intelligence and legacy. Historians describe the past with greater clarity, while scientists continue to push the limits of understanding. The global spread of literacy and the mass circulation of information through social networks has raised awareness of our history of innovation on planet Earth. The optimistic view is that *Homo sapiens* are becoming more enlightened in many ways, and new forms of collective intelligence are emerging as the networks connecting people with people, and people with machines, produce powerful new patterns of collective action across the globe. The pessimistic view is that these powerful new

patterns of collective action are neither intelligent nor adaptive, in the sense that people continue to make bad decisions.

One fact seems indisputable: people disagree with one another at the collective level. Anyone who dares speak of politics and religion is quickly made aware of this fact. In the greater scheme of things, consensus does not exist. While disagreements can drive innovation, they can also challenge solidarity. The nature of innovation is that it pushes outward in all directions, while its younger cousin, solidarity, seeks to pull everyone together again. Indeed, the historical development of science is marked by more diversity and specialism, and less transdisciplinary synthesis. In this context, the potential for greater collective intelligence, synthesis, and solidarity might be latent in the system, but the question is, *How can collective intelligence arise?*

The very notion of intelligence has an interesting history. A rather myopic view that emerged with the historical rise of psychological science (and the associated intelligence testing industry) was the view that intelligence is a property of individuals alone. In Western cultures that invested heavily in intelligence testing infrastructures this myopic view of intelligence reinforced individualist, competitive educational infrastructures that become increasingly dominant throughout the 20th century. Surprisingly, the notion that intelligence is a group-level or 'collective' phenomenon was rarely voiced by psychologists and education system designers. The dominant market forces and values shaping the behaviour of the industrialists of the 19th and early 20th century did nothing to help psychology and education break free of this view. Education system designers simply tracked the market, and the market demanded workers with particular skills – the skills needed to work in industry.

Fundamentally, the market illustrates and dictates what we value when it comes to educational development. In recent years, CEOs of multinational companies are crying out for graduates that possess an ability to work in a team structure, but this wasn't always the case. The dominant market forces of the industrialists of the 19th and early 20th century valued the efficiency of individual workers, and the capabilities of individuals acting alone. Dominated by the vision of Frederick Winslow Taylor and his industrial engineering principles of *scientific management*, industrial workers were largely seen as cogs in an efficient machine. Individuals worked on specific tasks, and were largely separated from one another in their work, with systems in place that maximized efficiency, compliance, and obedience above all else. Group intelligence was not valued, and was often feared. Fast forward to 2018 and we have barely escaped these primitive fears. As was the case throughout history, those who possess power fear the power of individuals and groups they cannot control. The consequence for modern *Homo sapiens* is that group intelligence has rarely been prioritized as a

valued educational outcome. Now, as the market pushes us toward global competitiveness and global cooperation, our uneducated minds struggle with the notion of collective intelligence. We may gravitate naturally and comfortably toward vague, mystical notions of community, god, and national spirit as this is what we have been taught to do. At the same time, given the complexity of the many societal challenges we face, and our growing recognition of the potential intellectual power of groups, it's unsurprising that the market is pushing us toward team structures. Still, our educational, societal, and work infrastructures contain shadows of Taylorism everywhere, and the focus on individuals continues to dominate our worldview.

As such, psychological science can hardly be blamed for its myopic view – the science was born under the disciplinary shadow of economics, and into a juvenile industrialised society. Furthermore, psychology adopts *a priori*, as its primary unit of analysis, a focus on the individual person. This *a priori* stance supports a particular focus of enquiry, and the development of bespoke methods for describing, explaining, and predicting individual human behaviour. At the same time, an anthropologist and a sociologist having dinner together, even if they agree on nothing else, might agree that the psychologists way of thinking about intelligence isn't altogether 'intelligent', given that, historically, *figuring things out* and adapting to the environment has been largely a group or collective effort for *Homo sapiens*. The fact that psychological science traditionally aligned itself with the strand of biological science focused on individual organisms doesn't uphold or make righteous the psychologists *a priori* stance *per se*, as biological science also has a strand of enquiry focused on groups and populations. The recent efforts by psychologists to create a 'new synthesis' that includes evolutionary psychology as a pillar of enquiry within their discipline doesn't necessarily balance the scales either, as the evolutionary psychology narrative on intelligence has tended to grapple with questions as to how evolution shaped individual intelligence, not collective intelligence.

But let's not be too precious about our various stances as we strut around the conference circuit and strike a pose for Vogue (while listening to Madonna on our headphones). There is nothing wrong with having an *a priori* stance. However, outside of one's *a priori* stance there is no requirement to grant any level of enquiry more status than another – cells, brains, individuals, small groups, and large groups are all open to selection. And when it comes to understanding intelligence, certainly, small group and large group dynamics are an important focus of enquiry. Indeed, aside from extreme scenarios where individuals were cast out by their tribe, roaming the forest alone, or placed into solitary confinement in prison cells, the history of *Homo sapiens* can reasonably be characterised as the history of groups *behaving together*, in more or less intelligent ways. While

group members may disagree in their value judgements as regards what is 'more' or 'less' intelligent, they continue to judge nonetheless.

As such, depending on where we live and how we have been educated, our culture of thought may force an imbalance in our stance on the subject of intelligence. But our current notions on the nature of intelligence should never be mistaken for the real history of intelligence-in-action. Still, if we're to redress the imbalance wrought by markets, education system designers, psychological science, and the intelligence testing industry, it will be important to complement our understanding of individual intelligence with a more advanced understanding of collective intelligence. As we move into this new phase of enquiry and understanding, we should continue to embrace psychological science and all the things that psychologists have figured out about intelligent action along the way. Thanks to experimental, cross-sectional and longitudinal research in psychology we certainly understand a great deal about the development of numeracy, literacy, and graphicacy in children and adults, how judgement and decision-making errors occur and are handled by individuals, the rate at which new knowledge is acquired and forgotten by individuals, how abstract ideas emerge from concrete representations, how environmental factors influence performance on intelligence tests, how intelligence test performance predicts academic, work, and health outcomes for individuals, and so on. The library is stacked with these books and I have read many of them.

However, as it stands, you won't find many books in the library with the words 'collective intelligence' in their title. We know much less about collective intelligence than we should. But this is about to change, and in his new book, *Big Mind: How Collective Intelligence can Change our World*, Geoff Mulgan wisely calls for the establishment of a new discipline of collective intelligence. Below I provide an overview of Mulgan's book and point to some collective intelligence work that is consistent with Mulgan's call to action.

Collective Intelligence in *Medias Res*

Mulgan defines collective intelligence (CI) as the capacity of groups to make good decisions using a combination of human and machine capabilities. Making 'good' decisions isn't easy, nor is working with groups, nor is integrating machine capabilities with human capabilities – so there are many challenges here. CI has been explored in a variety of different ways in the past, and Mulgan is aware that he is moving into occupied territory. For example, following Hegel's impulse, we have a variety of grand, idealistic narratives on CI, including Vladimir Verdandsky's account of the three stages of world development -- from the inanimate geosphere, to the biosphere of living things, to the Noosphere of collective thought and consciousness. Marshall McLuhan, Peter Russell, Gregory Stock,

and Pierre Levy have all written influential books pointing to the potential of new forms of CI. Howard Bloom extrapolated from bacterial and insect ecologies to describe how CI can emerge in human systems. Douglas Englebert and Eric Raymond extrapolated from the operational capacity of internet technology and the dynamics of open-source software developments to the potential augmentation of human collective capabilities. James Surowieki described research on the aggregation of individual decisions into collective decisions and argued a case for the ‘wisdom of crowds’. Mulgan also notes other traditions from sociology, anthropology, psychology, and multidisciplinary combinations, including the work of Harrison White and Mark Granovetter on the sociology of networks; Gabriel Tarde on group mind; and Mary Douglas on institutional arrangements shaping group cognition. Mulgan also points to Alex Pentland’s work on social physics, Dirk Helbing’s work on socio-technical systems, and Christopher Alexander’s work on pattern language, which was influential in the fields of urban design and computer software design. Furthermore, within management and organisational science, W. Edwards Deming, Chwen Sheu, Peter Senge, and James March address aspects of group work focused on variation, improvement, systems thinking, and decision-making; Ikujiro Nonaka writes about knowledge-creating organisations; and there are many scholars and practitioners working on management information systems, data management, mapping and mining; decision support systems, and the provision of services focused on organisational creativity, innovation, and change management. However, Mulgan’s primary point in signposting these contributions toward the end of his book is to remind us that, although many strands of enquiry, methods and findings are scattered across the landscape, a coherent discipline of CI has yet to be established. Mulgan does not attempt to integrate these strands *per se*.

Rather than begin with some effort to reverse engineer the history of CI and the rise to prominence of our new machine capabilities, Mulgan begins in *medias res*, in the middle of things, as they currently stand. Recent decades have seen many advances in tool and infrastructure design that have expanded the power of individuals and groups to think and act together in increasingly intelligent ways. Many of the examples of CI that Mulgan points to involve unique *assemblies* of specific human, machine, and organisational capabilities adapted to specific problems. For example, while roaming groups navigated by the stars for millennia, without even a rudimentary compass or map to guide their way, in recent years, our ability to navigate and explore new environments – on foot, via bicycle, or in our car – has been radically enhanced by Google Maps. As noted by Mulgan, the creation of Google Maps not only provided a step change in our collective capability to observe and navigate our environment, in developing the service Google exercised a unique form of collective intelligence design by assembling many distinct elements to make their service work. First,

they integrated various technology companies within their existing organisation and network, including Where 2 Technologies (providing zoomable, scrollable, searchable maps), Keyhole (providing geospatial visualisation software), ZipDash (supporting real-time traffic analysis via anonymous input from mobile phone users), and Vutool (providing street view imaging using a fleet of cars fitted with cameras). Google also used open source software methods to make it easier for other sites to integrate Google Maps, and they cut a deal with Apple to load Google Maps as a default app for all iPhone users; finally, they mobilised the collective intelligence of the public and provided them with ways and means to edit and add to maps using Google Map Maker.

Mulgan documents many other innovative assemblies that combine human and machine capabilities in efforts to support intelligent action on a large scale. Some assemblies leverage the network properties and tool affordances of the Web, for example, Wikipedia, which provides a web infrastructure to supports groups co-creating knowledge; Zooniverse, which supports collective engagement in research project work; Historypin, where people share photos and stories, telling the histories of their local communities; Thingiverse, a web community that includes more than 2 million users who have shared and remixed over 1 million 3D designs; Duolingo, which mobilized 150,000 people to test thousands of variants its web-based, automated language lessons, resulting in processes that allow people to learn a second language in around 34 hours.

Other assemblies use machine algorithms to perform specific types of intelligence work. For example, algorithms are increasingly used to guide investment decisions (e.g., VITAL is a data analysis system honoured with a seat on the board of a Hong Kong venture capital firm). Algorithms have been embedded in population profiling models that seek to predict the probability of prisoners reoffending. Algorithms are embedded in safety systems, including a system designed to predict the likelihood that any one of New York's 360,000 buildings will burn down in a fire. Mulgan notes how other assemblies combine observation and data, predictive models, and strategies to generate options and act in response to environmental and societal challenges. These include Hewlett-Packard's Central Nervous System for the Earth, Europe's Copernicus program, and The Planetary Skin set up by NASA and Cisco, each of which seeks to support adaptive global responses to environmental challenges (e.g., modelling and responding to extreme weather events, or shortages of energy, food, and water). Similar large-scale projects focus on population health, including MetaSub, which seeks to understand patterns of antimicrobial resistance by mapping the global urban microbial genome; and AIME, which seeks to track and predict outbreaks of Zika and dengue using a combination of machine and human intelligence.

Finally, says Mulgan, there are projects that have involved truly massive levels of combined human, machine, and organisational capacity. These include NASA's Apollo programme, which involved over 400,000 people, 20,000 firms and universities and the coordination of their combined resources and intelligence; and the Manhattan project, which employed 75,000 people, including many of the best scientists and engineers in the world, working in top secret silos across multiple isolated teams.

Although there are commonalities across these collective intelligence assemblies, in terms of the tools, talents, team processes and infrastructures used, these commonalities by no means reflect an orderly approach to collective intelligence design at a societal level. The assemblies of human and machine intelligence that have emerged in recent decades, says Mulgan, reflect as much the random happenstance of human ingenuity and innovation as they do any planned process of societal design. There is little by way of coherent vision, shared purpose, global knowledge exchange, and the design of any national and international collective intelligence infrastructures that would provide the hallmarks of a true discipline. And yet, a discipline of CI is sorely needed. In relation to the grand projects like The Manhattan Project and NASA's Apollo project, Mulgan recognises that they were conscious and diligent efforts to orchestrate human intelligence on a grand scale, but they had little or no influence on the dominant societal systems for CI, including how we organise our universities and our governments. Mulgan believes that many of our traditional, hierarchical institutional practices have remained somewhat fixed and frozen in time for centuries, and anyone walking into a University governance meeting or national parliamentary meeting today, 100 years ago, or even 200 years ago, would recognise the similarities. This lack of progress in the development of CI infrastructures is problematic.

When it comes to the design of new collective intelligence systems, says Mulgan, it's not a simple case of replacing the silence of traditional hierarchical systems (where no voice dared challenge the status quo) with systems where the chaos and confusing noise of undisciplined crowd behaviour inhibits the exercise of good judgement. Conscious orchestration of CI is needed. We need to design for CI in a disciplined and careful way. And if we fail to establish a discipline of CI – in schools, universities, governments, and in the interdependent exchanges between diverse groups at a societal level -- there is a risk that our biased, over-confident, manipulative, competitive, groupish tendencies will play out in more powerful and damaging ways across a range of technologically innovative, but poorly conceived, assemblies that dot the intelligence design landscape. For those who dream of a technological utopia, Mulgan argues against the view that we can rely on machine intelligence alone to address societal challenges. If we fail to establish a discipline of CI, there is a risk that CI will not keep pace with artificial intelligence (AI) -- decisions guided by the judgement

and wisdom of groups may fade into oblivion, and a dystopian scenario may arise whereby the wondrous pattern recognition, analytical and decision-making powers of machines sit amid a silent and mindless crowd who, sporadically, use ham-fisted and outdated methods for making decisions that matter most to *Homo sapiens*. The risk here is that *Homo sapiens* may become *Homo sine mente* if collective intelligence is not cultivated.

Much like brains rely on key structures and organisational dynamics to configure functional networks of neuronal activity and intelligent behaviour in context, so too do groups need structures and organisational dynamics to coordinate human and machine capabilities and CI activity in context. As Mulgan sees it, there is no need to place AI in opposition to CI. CI will no doubt act as an important counterpart and partner to the ever-strengthening discipline of artificial intelligence (AI). Rather than be afraid or deluded by dystopian fantasies, we simply need to think very carefully about our CI-AI designs. It is inevitable that our global intelligence infrastructures will evolve over the coming decades. If we invest in CI as a discipline, many insights, methods, tools and practices scattered across multiple fields can coalesce and co-evolve to enhance our conscious control over the systems we design and use. At the same time, a paradigm shift will be needed embed CI-AI within a broader transdisciplinary enterprise. Importantly, Mulgan provides a number of useful conceptual frames to support advances in the field.

In particular, Mulgan highlights a number of important dimensions of CI. These include:

- the *functional capabilities* of CI that are utilized and coordinated in context - observing, analysing, remembering, creating, empathising, and judging.
- *infrastructures* that are needed to support CI, including tools, methods, common standards and rules, and new institutions and networks.
- *principles* for organising CI work, in particular, autonomy, balance, focus, reflexivity, and integration for action.
- the three *learning loops* shaping CI applications, that is, using existing models to solve specific, pre-defined problems (loop 1); generating new models to support new learning (loop 2); and changing the way we think, including our epistemologies, ontologies, problem solving methods and tools (loop 3).

Mulgan also addresses the important theme of *cognitive economics*, or the ways in which we compute the costs and benefits associated with exercising CI. He considers the nature of collectives and the challenges of working with groups, and he highlights the importance of maintaining vigilance, self-suspicious and fighting the enemies of collective intelligence. Mulgan also attempts to

characterise some key stages in group problem solving, and he provides useful recommendations on how to run more effective group meetings. Finally, he highlights the challenges of promoting collective intelligence within Universities, Governments, and Society more generally. Overall, Mulgan provides us with a highly generative and powerful synthesis that will inspire current and future generations who seek to promote collective intelligence.

Functional capabilities

At the core of collective intelligence are a set of functional capabilities. Much like when we study individual intelligence, when we study collective intelligence, says Mulgan, it's important to identify the functional capabilities that help us solve problems and make decisions. These include the ability to construct a *model*, including a model of our 'world', or our current problem situation. Then we have a number of specific capabilities: to *observe*, *focus*, *analyse*, *remember*, *create*, *empathise*, and exercise *judgment* and *wisdom* in context. Groups or collectives can exercise these capabilities, and in so doing, they may perform at a higher level than an individual working alone. For example, a highly coordinated group can *observe* more than an individual alone can observe (e.g., when a group of people, as opposed to one person, diligently searches a supermarket or neighbourhood for a lost child). Similarly, a group of people can *remember* more than an individual working alone. For example, a group might coordinate their memories (e.g., semantic memory/knowledge) in response to a shared challenge (e.g., when a project team combines people with diverse knowledge to envision the design of a new city).

Beyond collective observation and collective memory, collective *analysis* and *judgement* can be challenging to coordinate, but even simple strategies can bear fruit. For example, a community leader might crowdsource solutions to a shared problem and thus has multiple people *analyse* a problem in parallel and propose independent solutions; a separate executive group might then *evaluate* the crowdsourced solutions and reward people who come up with the best solutions. In effort to exercise good judgement and make wise decisions, there is also potential for more closely coordinated group activity, for example, if deliberation spaces are designed where group members can share their knowledge, correct one another's logic, cross-check facts, elaborate richer understanding of a problem context, empathise across multiple dimensions of a problem field to ensure that all dimensions are made salient, balance one another's short-term and long-term preferences, and so on. Much of Mulgan's book focuses on design challenges related to these higher-order forms of collective intelligence, and situations where collective analysis, problem solving, judgement and decision making are critical (e.g., in the design of democratic systems, and in collective responses to societal problems).

As noted by Mulgan, each of the *functional capabilities*, or building blocks of CI, can also be enhanced by machines. For instance, we can use video surveillance technology and image analysis tools to help us *observe* and *analyse* our environment. We can use information systems to store data and uphold our collective *memory*. We can use computer software to analyse data (e.g., patterns of language used in lengthy documents or social media feeds; patterns in weather data that are used to predict future weather patterns; patterns in the stocks and flows shaping future business operations). However, Mulgan notes that machines struggle with creativity, judgement and wisdom, as the predominant machine designs, ancestors of Turing, are largely rule-based systems that struggle to generate new models and new ways of thinking. As Cathy O’Neil observed in her book, *Weapons of Math Destruction*, machine algorithms may reinforce and aggravate existing societal patterns, including patterns of inequality, and, indeed, may make them more difficult to change. Without a rule to follow, machines will fail to respond to the subtle, contextual features of a situation. Advocates of machine intelligence may argue that these are simply challenges that need to be overcome in the future, but, one way or another, the task of coordinating human intelligence and machine intelligence is a major design challenge going forward.

In a number of places throughout the book, Mulgan highlights the need for *balanced use* of different CI capabilities. For example, if a technology design group are fixated on *creativity* and innovation, but fail to cultivate a *memory* of past failures, they may not learn from their past and be doomed to repeat their history of past failures. Similarly, if a governance group are fixated on documenting, recording, and remembering every aspect of their past activity, the burden of their organisational *memory* may stifle *creativity* and every effort at innovation. Imbalance can manifest in any number of different ways, says Mulgan, and avoiding it requires an ability to gain perspective on the intelligence processes that play out in different situations, and modify CI processes if necessary. Invariably, there is scope to enhance both specific CI capabilities and balanced use of these CI capabilities.

Furthermore, Mulgan suggests we think about the *dimensionality of choices* in problem solving and decision-making contexts where CI plays out. The types of problems groups address can vary in their cognitive dimensionality (e.g., how many different models or ways of thinking are necessary to understand the choice), social dimensionality (e.g., how many individuals, groups, and organisations influence the choices being made; their degree of cohesion, conflict, cooperation), and temporal dimensionality (e.g., how long does it take before the group receives feedback to verify the value or validity of a choice). AI systems may work well in situations where there is well-ordered data and established rules and interpretative frames for making choices (e.g., making intelligent moves in a

game of chess). Societal problem solving, on the other hand, often involves groups working in situations where data is poorly structured, rules and interpretative frames are many and varied, and the dimensional complexity of the scenario is poorly understood and difficult to manage. As noted by Mulgan, many organisations attempt to address high-dimension problems with low-dimension tools. This tendency can be exacerbated in situations where an individual or group has the power to exclude others, and thus simply ignore dimensions of a problem situation that have been identified as relevant to decision-making. Exercising CI capabilities in a balanced way and in a way that does justice to the dimensionality of a problematic situation implies the need for infrastructures that support intelligence.

Infrastructures

Collective intelligence cannot be exercised in a vacuum – infrastructures are needed to support groups and teams that wish to coordinate their functional capabilities. These *infrastructures* include common standards and rules, tools and methods, and institutions and networks that support the intelligence of working groups. For example, without some common rules and standards in the use of language, it would be impossible for groups to communicate, learn together, coordinate their actions, and pass on their knowledge from one generation to the next. Mulgan provides a number of concrete examples. For instance, without the Trojan efforts of Richard Chenevix Trench and the massive team he assembled in the mid-19th century, we would not now have the Oxford English Dictionary as a shared resource that allows English speakers across the globe to coordinate common and complex meanings at a collective level. Similarly, consider the ways in which collective scientific and engineering endeavour depends on common rules and standards, including the many coding and measurement systems we use – the periodic table, biological taxonomies, the metric system, the Dewey Decimal System, medical diagnostic categorisation systems, semantic web tagging systems, and so on. While it is recognised that well-calibrated observation, sound measurement, and reliable coding systems are necessary for objective, standardized, coherent thinking and largescale thought and action in relation to many societal issues (e.g., using standard measures like GDP to track and compare societal progress across nations), Mulgan also notes how every ontology (i.e., every system for organising information) can result in a loss of information and a failure to establish shared understanding outside of the context where the ontology is used. The exercise of collective intelligence entails an ability to identify what matters in any given situation – to observe, analyse, judge and make decisions based on what is important. Caution needs to be exercised such that our ontologies don't blind us to what is important.

CI infrastructures also include *intelligent artefacts* – pen and paper, ruler and compass, excel sheet and smartphone application. Mulgan notes how intelligence is increasingly embedded in our physical and social environment – road markings, speed limit signs, and traffic lights keep us safe on the roads; prosthetic devices support our sight and mobility; sophisticated sensors and recording devices support observation and memory; and so on.

But CI requires more than intelligent artefacts – we also need to invest in human, social and organisational capital. Mulgan highlights some broad, positive trends over the past few hundred years: increased literacy rates; the rise in influence of governments, charities, universities; reduced crime levels, and increased levels of cooperation, trust, and sharing. Some areas of the world have developed more than others, and Mulgan notes how cities in particular have often been great crucibles of collective intelligence, with coffeehouses, clubs, societies and laboratories mobilizing intensive interaction between people, and with pockets of high investment of scarce resources catalysing creativity and collective innovation.

Social networks can also support knowledge and capacity exchanges that uphold collective intelligence. The growth in the number of learned societies around the world, says Mulgan, illustrates the trajectory of our cultural evolution in this regard (e.g., the Royal Society was established as one of the first learned societies in London in 1660; by 1880 there were 118 such learned societies). Today we have thousands of academic journals, and hundreds of thousands of annual academic conferences, seminars, and meetings, where people exchange knowledge and advance their disciplines.

In science, engineering, business, medicine, and government, a capacity to recruit, motivate, and manage teams, says Mulgan, is as important to the practice of CI as is the hardware and software that teams use. When it comes to managing groups and optimizing group size for CI work, Mulgan suggests that larger social networks are often good at gathering and organising knowledge, and they may also be good for argument and deliberation, but they work less well when it comes to decision-making and integration for action. Smaller groups and teams tend to dominate these roles. For example, projects like Wikipedia and Linux draw upon large networks of active contributors, but smaller, powerful groups of editors, curators, and guardians make many critical decisions that shape project work and key aspects of organisational strategy. When it comes to CI infrastructure design, says Mulgan, we need to design systems that support larger social networks and smaller problem solving, design, and executive teams to work together. Although Mulgan does not develop these design ideas or clarify specific design options, he does propose a number of organising principles that may be useful as part of the broader CI design process.

Organising principles

Mulgan argues that CI capabilities and infrastructures need to be assembled in ways that are fit for purpose in different contexts. A philosophy of creative design is central to Mulgan's mindset: while some *models of organisation* may generalise across contexts, creativity and care is always needed. Organising and operating a CI assembly always requires careful and bespoke designs. As such, rather than propose any specific models of organisation, Mulgan proposes five *principles* for CI assembly design: (1) autonomy, (2) balance, (3) focus, (4) reflexivity, and (5) integration for action. At this point in the book some readers may be disappointed that Mulgan focuses largely on principles, and bypasses the disciplinary value derived from describing, proposing, or reverse engineering specific organisational models (and specific methods) that groups might use to support their CI work. Nevertheless, by focusing on these five principles, Mulgan adds value and helps to frame important aspects of CI design thinking.

So how do successful CI assemblies work according to these five principles? According to Mulgan, successful CI assemblies:

- Create **autonomous** knowledge and informational commons, whereby the elements of CI are allowed free rein and are not subordinated to ego, hierarchy, ownership, assumption, fantasy or delusion. In contrast to the other four principles, Mulgan devotes a full chapter to the autonomy of intelligence. He notes that a group with more autonomous intelligence has a greater capacity for objectivity and self-correction, and is less likely to fall victim to errors of observation, memory, analysis, judgement and decision-making, including a tendency to confirm pre-existing beliefs, fixate on a narrow range of facts or perspectives, and become skewed by power and status, and so on. Infrastructures and institutions that support the autonomy of intelligence are already dotted across the landscape, says Mulgan, even if we don't always see them as such. For example, the use of the 'black box' in airplanes to record data and cockpit conversations supports important intelligence work conducted by the aviation industry and safety commissions as they seek to understand and avoid aviation errors, accidents, and disasters. Autonomous intelligence services are also used to support decision-making and the regulation of group behaviour, including the use of independent auditors to assess company accounts; a free media to investigate and report on current affairs; national and international think tanks that provide intelligence input to governments. The autonomy of collective intelligence implies that groups are open to making errors. Openness to error and potential failure is an important part of exploration and learning. Groups must also recognise the ambiguity of language and the limits of any

ontology they are using. This helps motivate groups to sustain an iterative process of communication and knowledge exchange as they work to establish shared understanding. More generally, the autonomy of intelligence implies that groups are free to exercise vigilance in efforts to maintain clarity, overcome biases, and establish shared understanding. An autonomous intelligence, free from constraint, is free to work with all ideas, and thus it can get carried away by imagination and fantasy, and potentially create models of the world that are divorced from reality. At the same time, only a fully autonomous intelligence allows for unbiased vigilance to manifest, says Mulgan, and only vigilance and the development of sound and self-correcting processes can sustain the collective intelligence of a group. Indeed, the next three principles that Mulgan advocates offer a definitive focus for the exercise of vigilance in a group.

- Successful CI assemblies, says Mulgan, achieve an appropriate **balance** between their functional capabilities. For example, as noted previously, CI assemblies might erroneously trade memory for creativity if they value innovation above all else. Similarly, CI assemblies can be rich in data, but poor in judgement and wisdom; or they can be rich in empathy and creativity, but poor in data. Mulgan doesn't provide much by way of additional design advice in this regard, but it is clear that an autonomous and vigilant intelligence must pay attention to any imbalance in the nature of its own process. In principle, the balanced and full use of all CI capabilities is needed to guide quality decision-making and collective action.
- Successful CI assemblies also achieve requisite **focus** on the task at hand, says Mulgan. Attending to what matters, ignoring distractions, and staying task focused can be challenging, particularly in situations where there is uncertainty as regards the nature of a problem a group is addressing. Mulgan doesn't elaborate on how different CI methods can be used to regulate focus. (A separate book or workbook on CI methods would be useful.)
- Successful CI assemblies **orchestrate** systematic reflection. They are reflexive in the sense that the intelligence process reflects upon itself, and is recursive and builds upon itself. The more explicit the data, predictions, processes, and learning outcomes, the easier it is to reflect upon and recursively evolve CI processes, and the more readily a group can exercise self-suspicion, critical thinking and iteratively redesign their approach.
- Finally, successful CI assemblies **integrate** for action. As Mulgan puts it, life depends on action, and part of the wisdom of intelligence is the ability to make the move from thinking to action. Although we must often broaden our view and embrace complexity to establish shared understanding and perspective, successful CI assemblies search for simplicity at the far side of complexity. Rather than become immobilized in a field of high-dimensional

choices, says Mulgan, successful CI assemblies have the capacity to move toward simplicity in the enactment of a particular decision, or a set of linked decisions.

Groups may fail to exercise these principles of CI in action. There are many reasons for this: groups may fail to even consider the principles in the first place; they may lose awareness of them as they work together; or regress to unprincipled modes of operation under conditions of stress or time pressure. Many influences disrupt and distort principled action, says Mulgan -- distractions, lies, rumours, heuristic thinking, opaque processes, trolling and spamming, to name a few. However, in an ideal scenario – where CI capabilities are present, CI infrastructures are in place, and a model of organisation that pays heed to the five principles of CI is used – a CI assembly should be capable of learning. Ideally, a CI assembly should also be capable of evolving its own learning process (e.g., by generating new models and new approaches to learning as needed). It is at this point that Mulgan introduces the idea of learning loops.

Learning Loops

According to Mulgan, CI assemblies and models of organisation need to allow for different levels of learning, each of which involves a different way of adapting to the environment. Mulgan distinguishes three learning loops. *First-loop learning* involves the application of existing models or methods to process data, analyse problems and make decisions. For example, in human-machine CI assemblies, machine algorithms can be used to support first-loop tasks if established rules and models of analysing data and arriving at a decision are available (e.g., analysing a play configuration and searching a decision space for the next best move in a game of chess). However, as Mulgan notes, computers may be powerful tools for playing chess, but they are not good at designing *new* games.

Specifically, computers struggle with the generative nature of *second-loop learning*. When there is a need to change the model that is used to support intelligent action in an uncertain environment, second-loop learning is needed. In this situation, learning must extend beyond existing rules to create new categories or relationships. An intelligent group can design new models when existing models no longer work. Second-loop learning requires an exploratory, reflective, generative stance, as new models are developed through a process of reflection.

Finally, a third loop of learning creates new ways of thinking, for example, by changing underlying ontologies, epistemologies, methods, philosophies of social organisation, systems of science and practice, and so on. In stable environments, where the rules of play are predictable, first loop learning may work well, says Mulgan, whereas in more changeable and unstable environments a

capacity to modify models and generate new ways of thinking is often necessary. This view resonates with the basic model of learning and development proposed by the influential psychologist, Jean Piaget. As Piaget described it, intellectual development is driven by the joint processes of organisation and adaptation – we organise information into mental models (or schemas) that support our ability to adapt to our environment. Our mental model of the world may work for a time to assimilate aspects of our reality (e.g., a child may get by for a time by calling every four-legged creature in the neighbourhood “Doggie”). But reality is always more complex than our simple mental models suggest, and when people are faced with new information (e.g., when the child’s parents say, “No, that’s a cat”), their mental model needs to accommodate this new information. Slowly but surely, a new level of organisation and adaptation is achieved -- a new and more complex model of reality emerges (e.g., eventually, the child distinguishes two species of quadruped, and can voice the correct label when these neighbourhood mammals are seen roaming around outside), and on the learning goes from there. Extending Piaget’s view of organisation and adaptation from individual intellectual development to collective intellectual development is useful in certain respects, but the challenge of developing and promoting CI remains. Indeed, building the discipline of CI requires significant investment in third loop learning, as we seek to generate useful ontologies, epistemologies, methods, systems of science and practice, and systems of social and political organisation that support the coordination of people and machines into the Big Mind that Mulgan hopes for. This is effortful, costly work and thus some consideration of the cognitive economics of CI is warranted.

Cognitive economics

Designing and implementing functional CI assemblies in response to societal challenges implies consideration of the *cognitive economics*, or costs and benefits, associated with exercising CI. Exercising CI may be beneficial, particularly as we work together to address societal challenges. But exercising CI can also be costly in terms of the time, energy, and funds needed for high quality CI work. The benefits need to outweigh the costs, and the value of CI work needs to be demonstrated to motivate continued engagement and investment of our scarce resources. Implicitly or explicitly, people think about the costs and benefits of participating in groups. Consider the challenge of starting a new group – a sports group, a hobby group, a lobby group – and the subsequent challenge of maintaining motivation and engagement, coordinating group member activity, and so on. The dynamics of cost/benefit analysis often sees people coming and going in different ways and it can be hard to keep a group together. Naturally, the same applies to CI work. Indeed, it may be harder to engage people in CI activities compared with other activities, and CI work does not necessarily come

naturally to us. Historically, we've seen massive investment and high levels of engagement with certain forms of group activity (e.g., sport, hobby, community, religious groups), but we haven't invested very much time, energy, or funds into CI group work. We also have a limited understanding of how successful CI work could be, and thus we have difficulty computing the potential benefits of CI. CI work is not universally valued in the same way as other types of group activity are valued. Unlike sport or hobby group activities, CI work may also lack key elements of enjoyment and play that motivates these groups to come together, and stay together.

We invariably recognise, on some level, that the operation of our intellectual capabilities has an associated cost. Individuals and groups may choose -- consciously or unconsciously -- to 'not' exercise their intellectual capabilities, particularly if the costs of exercising intelligence are perceived to be too high. The work involved may be judged to take too much time; bringing a group together may be judged to be too energy-consuming; acquiring the CI tools and methods, and hiring in methodological expertise, may be judged to be too expensive; and so on. As the discipline of CI emerges, resistance to the application of CI may be rule rather than the exception, particularly given that CI work may be more time consuming, energy-consuming and financially costly than other methods (e.g., decisions made by a leader or a machine algorithm). At the same time, as happened in our recent history of manufacturing, distributing, and embedding computers in everyday social and work life, embedding CI in our decision-making practice will become less costly if high quality infrastructures are developed to support the efficient design and delivery of CI activities in organisational contexts.

At a societal level, there are potentially massive costs associated with *not* exercising CI, says Mulgan. These include failures to innovate and capitalize on opportunities that can emerge from high quality CI work; the failure of groups to coordinate actions in response to societal and environmental problems; breakdowns in communication, trust, cohesion that result from minimal investment in CI work; repeated cycles of uncertainty, stress, and a collapse of adaptive resilience in groups who invest little effort in understanding the complexity of challenges they face; and so on. There's no escaping the reality - in order to make CI work, we need to invest time, energy and money to design, implement and experiment with new CI *infrastructures*. Crafting these new CI *infrastructures* will be important for the emergence of CI as a discipline. And if we discover CI methods that are working well, we need to invest considerable time, energy and money to *maintain* these infrastructures. Finally, if flexibility and innovation are part of our disciplinary efforts, we need to experiment, learn from experience and evolve our CI infrastructures, methods, and practices to continuously advance the discipline. A variety of mainstream disciplines need to converge in these efforts, says Mulgan,

including organisational studies, psychology, sociology, political science, philosophy, computer science and engineering, and economics.

Indeed, as part of the new discipline of CI, a sub-discipline of cognitive economics would be important to develop, much like the health sciences draw upon health economics to make decisions about how best to invest scarce resources to maintain the health of populations. As the discipline of CI stimulates the market for CI applications, a sub-discipline of cognitive economics would help us to work out the cost of CI tools, methods, manpower, organisational infrastructures, and the transaction costs of different CI activities. We could also address broader questions, says Mulgan: Is investment in CI more valuable in some environments than others (e.g., in situations where groups face more complex, fast-changing, and interdependent problem environments, as opposed to slow-changing, predictable environments)? What models should we use to compute the value of our organisational capital and the range of valued outcomes that derive from our cooperative activities? What investment strategy is needed throughout the process of formal education – from kindergarten to school to university – to support efficient and effective use of our CI infrastructures? What incentive systems are needed to motivate teams to engage with one another and exercise the very best of their collaborative intelligence? Outside of Mulgan's core narrative, there are other questions that will arise as the thinking of economists re-enters the political sphere with its new focus on collective intelligence. For example, to what extent might 'cognitive economics' displace the democratic requirement for political legitimacy and ethical evaluation of policies and courses of action? In an ideal world, it would not displace these things, but the world is not an ideal place, and it will be important to consider carefully the costs and benefits of different approaches to thinking.

More generally, while technology has advanced our powers of observation, memory, and analysis, as it stands, Mulgan sees a huge misallocation of brain-power and machine intelligence, with some of the most important fields, including politics and education, currently possessing little by way of CI capabilities and infrastructures, and thus remain locked in action patterns that limit their effectiveness. But action patterns are open to change, and now is a good time for cultural evolution.

CI in Action - meetings

As noted by Mulgan, collective intelligence efforts play out most frequently in meetings – workplace meetings, town hall meetings, executive committee meetings, conference meetings, and so on. Traditionally, there is a common format to many of these meetings: a small group works through an agenda, they make some effort to take turns speaking, and they deliberate together in an effort to

make decisions. Technology innovations have not radically altered traditional meeting formats, and meetings are often perceived by groups as unproductive, frustrating, and a waste of time. New technologies -- google hangouts, webinars, fancy powerpoint slides, data visualisations, etc. -- may do little to offset the frustration that people experience in meetings, particularly if fundamental aspects of group dynamics (e.g., extraverts dominating the airwaves) continue to disrupt the productivity and collective intelligence of a group. Mulgan notes that new meeting formats (e.g., World Cafés, Flipped Learning Conferences, and Holocracy) may at first glance appear more egalitarian, open and dynamic, offering group members autonomy over the flow of ideas and the conversations they contribute to. But these new meeting methods can also be overly vague, difficult to organise, dominated by extroverts, and unsuitable for sustained problem solving, says Mulgan.

At the same time, for groups working to address complex problems in rapidly changing environments, regular meetings allowing for knowledge exchange and deliberation are often necessary, and these meetings are often important in efforts to maintain smooth social relations in the group. To work well, says Mulgan, meetings need to counter our tendencies to (1) favour social harmony over a desire to share novel or discomforting information, (2) become attached to ideas in a way that makes it hard for us to see their flaws, (3) defer to authority (thus sacrificing our intellectual autonomy), and (4) defaulting to equality, in the sense that the value of independent contributions is not evaluated *per se*, with each idea treated as equal.

In order to support collective intelligence during our meetings Mulgan proposes the following:

- 1) Make sure all participants understand the purposes, structures, and content of the meeting; share an agenda and any information in advance of the meeting such that everyone is up to speed and time is not wasted.
- 2) Make sure the meeting is facilitated well, such that group members stay focused on goals, take turns speaking, pause for reflection, and so on.
- 3) Encourage group members to articulate and interrogate arguments; allocate people roles to interrogate arguments; and provide a structure and incentives for dialogue and argumentation as needed.
- 4) Use multiple communication formats to aid understanding and learning (e.g., combining written, visual, and verbal communication formats), and keep in mind simple communication rules, such as: “no numbers without a story, and no story without numbers, or no facts without a model, and no model without facts.” (p. 136).
- 5) Rein in extroverts and opinionated and powerful people if they seek to dominate meetings, and seek to promote reciprocity and social perceptiveness of group members (e.g., by

creating groups that include both women and men, which may enhance reciprocity and the average social perceptiveness of groups, given that women score higher than men on measures of social perceptiveness).

- 6) Design the physical space in a way that supports collective intelligence work, for example, using furniture and seating arrangements that allow for productive dialogue and exchange, and by providing sufficient natural light and space to move, etc.
- 7) Design for effective division of labour, for example, by using methods that allow group members to adopt different roles, and to switch between these roles as needed. For example, Mulgan mentions Edward de Bono's method, which allows people to wear different 'thinking hats', which prime them to either engage in critical thinking (black hat), focus on feelings and intuitions (red hat), manage thinking processes for the group (blue hat), promote creativity (green hat), or engage in optimistic lines of thinking (yellow hat). Mulgan also mentions Kantor's method, where groups members adopt one of four roles: "movers, who initiate ideas and offer direction; followers, who complete what is said, help others clarify their thoughts, and support what is happening; opposers, who challenge what is being said and question its validity; and bystanders, who notice what is going on and provide perspective on what is happening, offering a set of actions people can take while in a conversation" (p. 138).
- 8) Seek a balance between the complexity or breadth of the topic being addressed at the meeting, the knowledge and experience needed to address the topic, the number of participants needed, time needed, and the degree of shared common grounding, language, or understanding needed to uphold collective intelligence. Although no empirical evidence is provided, Mulgan suggests the following formula: $\text{meeting quality} = [\text{time} \times \text{common grounding} \times \text{relevant knowledge and experience}] / [\text{numbers} \times \text{topic breadth}]$
- 9) Make meetings visibly cumulative, for example, using feedback forms and reviews that link consecutive meetings over time; make use of dashboards that track progress; and leverage social network analysis tools to track how people interact after meetings.
- 10) Do not allow a culture of wasteful meetings to result in undue frustration and boredom. Instead, develop a culture where it is acceptable to cancel meetings if they are not needed, shorten meeting duration to align with the number and seriousness of issues being addressed, and consult with group members in advance on how long a meeting should be or whether or not it's needed. If the meeting is needed, ensure the culture of the organisation upholds key principles of CI, specifically, by (1) ensuring the autonomy of intelligence, deliberation, and problem solving efforts, (2) ensuring balanced use of different functional

capabilities during the meeting, (3) maintaining focus on the goals of the meeting, distinguishing what is relevant from what is not relevant, (4) reflecting on the intellectual contributions throughout the meeting and considering if new thinking methods are needed, and (5) allow the facilitator or chair of the meeting to work with the group to integrate for action, ensuring some effort is made to move from the complexity of the understanding developed by the group to the simplicity of a good decision.

Solving complex problems on the scale of cities and beyond

Beyond small group meetings, CI in action can and should manifest on a larger scale, for example, in the way we run our cities. Some examples of third-loop learning, says Mulgan, have arisen in cities that have invented new ways of thinking about their past and current condition and future prospects. Mulgan notes, for example, how 19th century pioneers in statistics developed methods to document and analyse the prevalence of poverty and diseases in cities, and how the planning professions emerged to prominence to support the design of new transport systems and parks and recreational facilities. Engineering, medicine, and social science methods were co-opted and integrated in innovative ways to shape infrastructure design, maintenance, and growth of major cities. The transition from the late 20th to early 21st century saw the emergence of new multistakeholder partnerships (e.g., the Basque region's Bilbao Metropoli-30, and the London Collaborative) that aimed to catalyse collective intelligence, resolve societal problems and shape the future design of cities and regions. Human intelligence was increasingly supported by machine intelligence. For example, Rio de Janeiro's famous control room allows city staff to analyse real-time data on air quality, noise levels, traffic congestion, dengue fever outbreaks, landslide risk zones, and thousands of GPS-tracked buses and ambulances. Mulgan notes how city administrators across the world now rely upon data analytics offices and services that analyse and package data in an effort to support ongoing problem solving and decision-making activities. The smart cities movement continues to push the limits of intelligence design in an effort to enhance the operational efficiency of cities and the well-being of city dwellers.

To illustrate the process of group problem solving, or collective intelligence at the level of cities, Mulgan draws upon his experience working with a team in central government to reduce homelessness across cities in the UK. Although his writing in this chapter is somewhat vague, with scarce methodological detail provided (see Chapter 12), Mulgan broadly characterises a number of steps in the process of problem solving: (1) circling and digging, (2) widening, (3) narrowing, and (4) iterating. *Circling and digging* involves an effort on the part of the group to describe the problem they are addressing, its dimensions and scale, how difficult it is and how many organisations need to

work together to address the problem, what established facts and data are available and how relevant and useful they are. Although no specific problem definition methods or systems thinking methods are mentioned, Mulgan broadly suggests that analytical tools involving simulations, models, and scenarios and the import of research findings from various disciplines (e.g., psychology, economics and markets, and political science) may be useful during the *circling and digging* stage of problem solving. Next, *widening*, involves the generation and prompting of a set of options, solutions, or tools that may help to address the problem that has come to be understood during the circling and digging phase. Some solutions may be novel, but others can be borrowed from elsewhere (e.g., when working to reduce the number of people sleeping on the streets, Mulgan's team borrowed ideas from work in the areas of case management, pooled budgets, preventative interventions, and city-level partnerships). When it comes to proposing solutions, Mulgan mentions Linus Pauling's edict: the best way to have good ideas is to have lots of ideas and throw away the bad ones.

Narrowing is the next step and it involves selecting solutions from the full set of options available, using specific criteria such as how feasible, costly, and sustainable solutions are and how strong the evidence is in support of their potential efficacy. Narrowing involves probability judgements and the relative weighting and aggregation of factors that have a bearing on the decision to select or reject a specific option. As Mulgan says, the more uncertainty and novelty there is in a problem situation, the more difficult narrowing becomes – and here a group risks relying on intuition, analogy to seemingly related problems, or they may simply defer to authority opinion. In a situation where there are viable solutions to a problem, and enough time and money available, it can be useful to experiment with variants and see which solutions work best.

Finally, says Mulgan, *iterating* involves recognition that public problems like homelessness are not like mathematical problems – any solutions selected for implementation are likely to be imperfect and will need to be carefully evaluated, modified, and iteratively redesigned and potentially replaced with new solutions (e.g., if the solutions implemented generate new problems that the CI group needs to respond to in the future).

A truly smart city will not only have the infrastructure and resources to support teams working continuously to resolve public problems -- circling and digging, widening, narrowing, and iterating as they work -- a truly smart city will also have ways of handling problems at different levels in a triggered hierarchy. Central to Mulgan's view more generally is the idea that successful CI assemblies – within organisations, cities, and even nations – have the capacity and resources to move from first, through second to third loop learning in triggered hierarchy as needed. Consistent

with first-loop learning, for example, a smart city will have many automated and standardized ways of dealing with simple, predictable and repeated tasks involving law and order, tax collection, education and health care delivery, and so on. But a smart city will also have higher tier authorities that support second-loop learning, specifically, when new models, rules or procedures need to be developed (e.g., when the health and well-being of commuters is compromised and new public transport models need to be created). A smart city will also have a capacity for third-loop learning if new approaches to thinking are needed (e.g., if the methodology used to design transport systems needs to be overhauled and a new paradigm developed).

One important paradigm shift that gathered pace in recent decades is the use of increasingly sophisticated public-private collaboratives. Mulgan mentions Cincinnati's living laboratory established in 1916 (following the model of collaboration developed by John Dewey), and the Peckham Experiment in the 1930s in south London, both of which involved teams working to innovate solutions to public problems. Mulgan notes how both were eventually crushed by political opposition, and he comments on the fragility of CI infrastructures in the face of mainstream political power. But new innovations continue to emerge, many of which leverage internet technologies. These include the US government's challenge.gov, an open innovation hub which was visited by 3.5 million citizens from 2005 to 2010 and hosted 400 challenges. Mulgan describes his experience working with the UK government on strategic audits across multiple policy issues, and he notes that honest, transparent, reflexive modes of operation are still rare in government. While many aspects of government are potentially open to inspection and improvement, the norm of openness will need to become more prevalent as part of our mainstream political philosophy. Further experimentation will also be needed to effectively build reflexivity and the three loops of learning into institutional practices.

Democratic assembly

Although democracy developed in part as a way to protect the public from oppression, says Mulgan, offering a means for citizens to influence policy via engagement with their elected representatives, modern democracy is flawed and considerable work is needed to enhance the collective intelligence of political parties and the people they represent. With the rise of political science and the spread of media, we are increasingly aware of the many flaws in our characteristic democratic systems. Mulgan highlights many such flaws: voters choose their representatives and parties out of loyalty rather than by reflecting on their policies; public attitudes toward political representatives and their policies are distorted by the media; our political representatives can be misguided, ignorant, and corrupt; and political parties often fail to think clearly and fail to use valid methods when

formulating policies. As Mulgan notes, “most of the machinery of contemporary political parties is devoted to campaigning not thinking...” (p. 190-91). Significant time and mental energy is devoted to political manoeuvrings and efforts to hold on to power. These political dynamics are hard to change.

As noted by Mulgan, utopian visions of direct democracy – aided by new online petition tools such as change.org and avaaz.org – fail to recognise that presenting citizens with simple binary choices (e.g., vote ‘yes’ or ‘no’ on a proposed policy) is of little benefit if citizens are not educated in relation to the issues or if they have no opportunity to engage in dialogue with others in advance of voting. Popular assemblies can be swayed by eloquence, charisma, and emotion, and people may be unduly persuaded by the opinions of others in their social media bubble. Those with power can sabotage legitimate institutional processes and use corruption to gain the support of the mob. More than simply choosing between alternatives, says Mulgan, a robust deliberative democracy depends on a strong civil society, a culture of trust and mutual respect, an independent media, an infrastructure that support knowledge generation, and reliable and valid methods that tap into the wisdom of citizens. Mulgan notes that more public participation alone does not necessarily lead to better decisions and it can lead to poor decisions (e.g., the anti-flouridization votes and plebiscites on taxation). What is needed, says Mulgan, are processes that “amplify thoughtful inputs and constrain the less constructive ones” (p. 184), and institutions that serve truth and uphold the autonomy of intelligence. Having said that, Mulgan doesn’t clarify what he means precisely by ‘thoughtful’ and ‘less constructive’.

Mulgan makes reference to a cluster of organisations and tools that he helped design in the UK, including the Alliance for Useful Evidence, the What Works centres designed to deliver facts and evidence into the system of government decision making; and the D-CENT tools developed by NESTA, which provide the public with an opportunity to propose issues, suggest policies, comment and vote. Other countries have similar infrastructures supporting the creative and critical thinking of citizens, for example, the Open Ministry in Finland, which allows the public to propose legislation and comment on ideas. Initiatives in Paris and Reykjavik go a step further and have citizens promote and rank initiatives or options for public spending, thus allowing for participatory budgeting at the local city level. Spain’s Podemos party used similar methods to shape policy in a number of cities. However, as noted by Mulgan: “None of these models is yet mature. The newer parties are still uneasy about taking on the responsibilities of power. Too many of their tools are more expressive than epistemic, and the more direct forms of democracy are still in competition with the older ones.” (p. 186). Notably, while online platforms (e.g., Loomio, Your Priorities, and DemocracyOS)

that allow for the anonymous input of ideas can be useful for identifying issues, generating options and commenting on options, Mulgan observes that when it comes to decision-making accountability matters and it is less acceptable for citizens to be 'anonymous'. However, there are many genuine design challenges to consider here, and Mulgan notes that a move from large- to small-group decision making may be necessary for difficult decision that involve consideration of trade-offs. Iterative and cyclical patterns of coordination between larger and smaller groups may provide part of the solution. For example, after small groups of representative stakeholders and content experts have considered trade-offs and proposed policy actions, with it comes to scrutinizing actions that emerged from decisions, a larger group involving civil society, universities, and independent organisation can again make valuable contributions. Clearly, says Mulgan, separating out steps in the democratic decision-making process is important -- framing questions, identifying issues amenable to action, generating options, scrutinizing options, deciding what to do, and scrutinizing what has been done. Also, maintaining independent roles for intelligence groups and policy decision-making groups is also valuable, say Mulgan. For instance, the International Panel on Climate Change analyses facts and scrutinizes proposed solutions; politicians decide between alternatives; and the action of government is then evaluated by universities, civil society, and independent organisations.

At the same time, the risks of using online platforms are understood by Mulgan, and it is clear that we have a lot to learn about the optimal design of these platforms. For example, for issues involving deeply held beliefs or values (e.g., gay marriage), online debate may aggravate negative feeling between opposing groups, and result in more polarisation and poor quality deliberation. Other issues (e.g., monetary policy; or policies for responding to threatening epidemics) may require high levels of technical knowledge and be unsuitable for public debate. Quality public debate also requires careful orchestration and moderation, balanced input from introverts and extroverts, those more or less well connected, and people who are capable of responding to and synthesising inputs from the public (e.g., a trusted public administrator). Quality public deliberation also implies the design of spaces where citizens are educated in relation to the issues, and where legitimate decisions are seen to emerge from well-informed deliberation. A well-designed democratic system, says Mulgan, needs to support all core functional capabilities of CI – observing reality clearly, remembering well, focusing on the problems that matter, reasoning clearly, being creative when exploring possibilities, and making wise judgements using sound methods.

In order for democracy to develop, a collaborative ecosystem focused on quality CI work is needed -- parliaments in collaboration with citizens, universities, media, think tanks, independent organisations, technology companies, and so on. Training and preparation for public administration

is also important, says Mulgan, although this training and preparation is rarely in political systems across the world. China is one of the few exceptions. Mulgan believes that political leaders “should be assessed for their readiness for their jobs, should be trained to fill the gaps, and should learn systematically on the job. We should want systems that are sufficiently open that incumbents can be sacked, but not so fluid that they are run by amateurs, which implies a bigger role for education on the job.” (p. 191).

Finally, says Mulgan, we need to think across different scales, and thus the democratic assemblies designed for running cities will be different from the assemblies designed for running nations. While crowds can help with many tasks, Mulgan believes they are poorly suited to the job of the design of new institutions, the crafting of strategy, or developing coherent programmes by combining discrete policies. Furthermore, a democratic system that upholds quality CI work cannot simply be deduced from principles alone – we need to attend to the details of the design and evolve our designs in light of experience and feedback. A creative and experimental mindset is needed for the design of effective democratic systems, as the emergence of good design is a slow, iterative process.

Collective wisdom and progress in consciousness

While some commentators optimistically point to a progressive trend in the evolution of intelligence, consciousness, and organisational and political practices (e.g., consider the work of Clare Graves, Ken Wilber, Ray Kurzweil, Federic Laloux, and Robert Wright), Mulgan considers this narrative as overly confident and linear and he warns that we cannot assume “an onward march toward more connected and thoughtful societies” (p. 217). For instance, while embedded forms of machine intelligence are predicted to transform whole sectors from transport to law and medicine, Mulgan notes that the societal changes will be dialectical and nonlinear as new demands will emerge as many current demands are increasingly automated and fulfilled by machines. As such demand may increase for coaching, supervision, craft work, and face-to-face care services, all of which may show a relative price increase. While technologists and futurists may assume that “the public is dumb and passive ... Two hundred years of technological revolution should have taught us that technological determinism is always misleading – mainly because people have brains as well as interests. People campaign, lobby, argue, and organise....they become agents rather than victims.” (p. 219). The dynamics of human politics invariably drives a dialectic change process and transforms simple-minded utopian or dystopian futures into mere speculation in retrospect. As part of a political process, issues of machine capability, human reengineering, and societal change will be hammered out one way or another, but it’s a mistake to assume we can fully predict the outcomes in advance.

If sufficient resources are invested and the hard work of CI design is embraced by politicians and citizens, collective intelligence assemblies could evolve in a progressive way – with flexible configurations designed on smaller and larger scales that combine all the key elements of intelligence. However, as it stands CI assemblies are rare, says Mulgan, and only a handful of assemblies like MetaSub and the Cancer Registry combine multiple functional capabilities. According to Mulgan, most CI assemblies have a narrow focus, including just one or two aspects of intelligence (e.g., Google Maps supports observation; Wikipedia supports interpretation and memory). Few have a reliable funding base, or access to a network of skilled professionals, and there is no established discipline of CI to support the design, implementation, and evaluation of new assemblies. Politically, we need to recognise the importance of collective brainpower – we need more “collective intelligence about collective intelligence” (p. 221). Mulgan tells us that there is a misallocation of brainpower to certain sectors – in particular, the military, finance and banking -- with much less public brainpower devoted to education, food, and energy. Brainpower in the private sector tends to follow the money trail, says Mulgan, with major investment seen in high profit, competitive areas such pharmaceuticals and computing.

Between the conservative fear of change and the liberal enthusiasm for change, a more mature politics could “fight for wider access to the tools of intelligence and a better allocation of those resources to the things that matter.” (p. 223). In this scenario, the intelligence, productivity, and democratic empowerment of individuals and groups could be enhanced and we could achieve “the greatest agency for the greatest number.” (p. 223). We need to figure out ways to work across global and local territories, recruiting whatever knowledge and capacity we can in efforts to respond to societal challenges. While science and technology advances our knowledge and capacity to control more and more aspects of our world, we need to integrate our global knowledge and capacity with our local knowledge and capacity. Advancing our wisdom and our ability to make better decisions collectively involves making the local context and the ethics of decisions visible; it involves thinking about the short- and long-term implications of decisions; it involves perspective-taking and empathising with others; it involves extending beyond the boundaries of the self, understanding one’s position in relation to the collective, and maintaining an openness to the broader world. We can imagine all sorts of futures, and our imagination can prompt all sorts of design efforts and experiments. At one level, we might imagine a world where few forms of collective intelligence have evolved – where our minds merge with machine intelligence, where the illusion of self is transcended, where our conversations in the spaces between one another are made visible and connected. While the prospect of this new Big Mind might terrify some, says Mulgan, only a culture and consciousness more advanced than ours can evaluate the unknown possibilities

and, one way or another, it is likely that we will continue to experiment and debate and transform our intelligence landscape. Rather than be terrified, Mulgan suggests we might adopt a more poetical stance: “Better to think, with William Butler Yeats, that “the world is full of magic things, patiently waiting for our senses to grow sharper.” (p. 228)

Beyond the magical potential, Mulgan recognises the hard work of CI discipline development will require a critical mass of engagement from established disciplines. But we have much to work with -- methods to analyse the costs and benefits of different forms of cognition (economics), tools to analyse group dynamics and social networks (psychology and sociology), approaches to understanding cultures (anthropology), and methods for understanding and enhancing pattern recognition and learning (computer science). A new synthesis of theory, research methods, and applied design work is latent in the system. It would help, says Mulgan, if Universities could reinvent themselves by, first, learning more about themselves as centres of learning. University governance models and the traditional modular model of teaching delivery tends to reinforce disciplinary divisions and power hierarchies that inhibit innovation and transdisciplinary CI work. Orchestrated experimentation and systematic evaluation of new models of teaching and learning are needed, but this should involve a transdisciplinary focus on the university itself rather than a focus on innovation within specific disciplines in the university. While internet innovations like the massive open online courses (MOOCs) appeared to break new ground, benefits have been limited in practice and these services have done nothing to promote collective intelligence *per se*. Rather than organise universities around separate disciplines and separate organised bodies of knowledge, an alternative involves organised learning (both theoretical and applied) around specific problems that call upon teamwork across multiple disciplines. Mulgan mentions Stanford University, Imperial College London, and Tsinghua University as leading the way in this regard, although he doesn't specify what they are doing that is altogether unique and distinct from many other Universities. (I can hear the creaking of chairs as college deans around the world shift and grumble.) The university, says Mulgan, could place a stronger emphasis on the orchestration rather than dissemination of knowledge – orchestrating, experimenting, dissemination, and evolving the discipline of collective intelligence by investing in processes, institutions, people, funding models that drive innovation.

This aligns with the work we have been doing at the National University of Ireland, Galway. In 2018 we established a new Collective Intelligence Network (CIN), and the associated Collective Intelligence Network Support Unit (CINSU), which provides facilitation support for teams and groups working to address complex issues across a variety of organizational and societal contexts. CINSU includes facilitators from multiple disciplines and backgrounds who strive to deliver high quality

collective intelligence facilitation. CINSU members use John Warfield's methodology, Interactive Management, combined with other collective intelligence methods (e.g., scenario-based design) to maximize team intelligence and collective action potential.

John Warfield, past president for the International Society for Systems Sciences, was a visionary thinker. He had a vision for applied systems science. He developed his Interactive Management methodology to support the work of applied systems science. In Warfield's model, the process of CI work involves the application of a set of methods and tools that helps groups to develop outcomes that integrate contributions from individuals with diverse views, backgrounds, and perspectives. Warfield's CI method can help to support high quality interdisciplinary work. We have found the method to be very useful in the design phase of major EU projects, including [Route-To-PA](#), [Q-Tales](#), [Sea for Society](#), [OpenGovIntelligence](#), and [INCLUSILVER](#).

Warfield argued that when (1) a *team* of people come together to focus on (2) a complex *issue*, they need (3) a *methodology* that helps them achieve an adequate synthesis of knowledge and perspectives that supports collective understanding and action in response to the issue. Warfield highlighted the need to partition the team into three sub-groups:

- Stakeholders – the people who have a stake in the issue being considered.
- Content specialists – the people who have specialized knowledge that is relevant to an issue under consideration.
- Structural modelers – the people whose task it is to structure the issue being considered.

While stakeholders and content specialists communicate the knowledge essential for understanding the issue or problem the team is addressing, the structural modelers facilitate the team in structuring their knowledge using specific facilitation strategies and tools.

Central to Warfield's CI methodology are a number of steps. First, a group comes together to generate, clarify, and select ideas relevant to the problem they are addressing (see step 1 & 2 in Figure 1). Using matrix structuring software, *Interpretative Structural Modelling* (ISM), key problem issues are systematically structured in pairs and the same question is asked of each in turn (e.g., "Does A influence B?"). After all the critical issues have been analysed in this way, the matrix structuring software generates a graphical problem structure (or *problematique*) showing how the issues are interrelated. The *problematique* can be viewed and printed for discussion (see steps 3 & 4 in Figure 1). The *problematique* becomes the launch pad for planning solutions to problems within the problem field. Warfield's CI method has been applied in many different situations to accomplish many different goals, including mediating peacebuilding in protracted conflicts (Broome, 2006;

2017), improving tribal governance process in Native American communities (Broome, 1995a, 1995b; Broome & Christakis, 1988; Broome & Cromer, 1991), developing a national wellbeing measurement framework (Hogan et al., 2015), and mobilising communities across Europe in response to marine sustainability challenges (Domegan et al., 2016).

Figure 1: A simple visual description of some of the key steps in the CI methodology

Within the local University context where CINSU was established in 2018, project work is focused on:

- Supporting University Governors with ongoing Risk Management Project Work
- Supporting Well-Being project work for NUI, Galway, Hardiman Library Staff
- Supporting Well-Being project Work for the NUI, Galway, LGBT+ Staff and Student Network

Also, in a recent application of Warfield's method, the CIN met to consider the challenge of integrating content expertise and methodological expertise in a team-based setting in efforts to address social issues. With expertise in education, business, psychology, sociology and politics, physics, and marine science, the CIN represents a diverse group of academics and practitioners with experience working across a broad variety of societal issues.

The group generated ideas in response to the following trigger question:

In the context of the design and implementation of solutions for complex social issues, what are challenges to integrating content expertise and methodological expertise in a team-based setting?

A systems model depicting connections among the set of ideas was also developed using Warfield's method (see Figure 2). The model is to be read from left to right, with arrows connecting challenges indicating 'significantly aggravates'. The model is easy to understand. For instance, to the far left of the structure is the *failure to correctly identify and include all of the key stakeholders*. During the structuring process, participants highlighted how this failure significantly aggravates *Failure to systematically plan and prepare for integration of content and methodological expertise*, which in turn significantly aggravate *Shortage of platforms, methods, and procedures enabling different views to be heard*. Collectively, these challenges were seen to aggravate challenges in the cycle to the far right of the structure - inertia, lack of openness, failure to question, excessive specialisation, and unequally empowered stakeholders.

In conclusion, while it has been noted that genuine human progress is most often collaborative and the result of people utilizing their diverse knowledge and skills in environments that support effective communication and teamwork, consistent with Mulgan's view, the CIN at the National University of Ireland, Galway have identified a need to better design and support teams in efforts to resolve societal problems into the future. Building upon the work of John Warfield, CIN members have previously argued that a pragmatic approach is needed to support teamwork and collective intelligence. John Dewey famously argued that a democratic society requires an educated population who can collaborate, deliberate and learn together. However, if skilled and knowledgeable transdisciplinary teams are to become the leading edge of societal problem efforts into the future, a major challenge for the current generation is to design a closely coupled educational and political infrastructure to support teamwork skill development and the slow, empowering path of participatory democracy. Beyond disciplinary and political divides, developing new strategies and behaviours that support teamwork and cooperation will be essential for the future survival, adaptation, and flourishing of *Homo sapiens*.

Figure 2: Structural Model of Interdependencies Between Challenges

The Arrow Should be Interpreted as:
“Significantly Aggravates”

- Makes it more difficult to resolve
- Increases the Severity of
 - Makes Worse
 - Exacerbates

The letter-number combination after each item represents the category from the list below and the number-order from Table 1.

- A. Stakeholders
- B. Group Dynamics
- C. Planning
- D. Resources
- E. Defining Goals

- F. Engagement
- G. Communication Issues
- H. Resistance and Fears
- I. Collaborative Methodology
- J. Epistemic Heterogeneity

Featured Book:

Mulgan, G. (2018). *Big Mind. How Collective Intelligence Can Change Our World*. Princeton University Press: Princeton, NJ.

Relevant Readings:

Broome, B. (1995b). The role of facilitated group process in community-based planning and design: Promoting greater participation in Comanche tribal governance. In L. R. Frey (Ed.), *Innovations in group facilitation: Applications in natural settings* (pp. 27-52). Cresskill, NJ: Hampton Press.

Broome, B. J. (1995a). Collective design of the future: Structural analysis of tribal vision statements. *American Indian Quarterly*, 19, 205-228.

Broome, B. J. (2006). Applications of Interactive Design Methodologies in Protracted Conflict Situations, in Lawrence Frey (Ed.), *Facilitating group communication in context: Innovations and applications with natural groups*, pp. 125-154, Hampton Press.

Broome, B. J., & Christakis, A. N. (1988). A culturally-sensitive approach to tribal governance issue management. *International Journal of Intercultural Relations*, 12, 107-123.

Broome, B. J., & Cromer, I. L. (1991). Strategic planning for tribal economic development: A culturally appropriate model for consensus building. *International Journal of Conflict Management*, 2, 217-234.

Broome, B. J., & Fulbright, L. (1995). A multi-stage influence model of barriers to group problem solving. *Small Group Research*, 26, 25-55.

Broome, B.J. (2017). *Mediating Peacebuilding in Protracted Conflicts: An Interactive Design Framework*, in A. Georgakopoulos (Ed), *The Mediation Handbook: Research, Theory, and Practice* (pp. 379-387), Routledge.

Domegan, C., McHugh, P., Devaney, M., Duane, S., Hogan, M., Broome, B. J., et al. (2016). Systems-thinking social marketing: conceptual extensions and empirical investigations. *Journal of Marketing Management*, 1-22.

Forsyth, D. R. (2014). *Group dynamics* (6th ed.). Belmont, CA: Wadsworth Cengage Learning.

Groarke, J. M., & Hogan, M. J. (2016). Enhancing wellbeing: An emerging model of the adaptive functions of music listening. *Psychology of Music*, 44(4), 769-791.

Hogan, M. , Broome, B. , Harney, O. , Noone, C. , Dromgool-Regan, C. , Hall, T. , Hayden, S. , O'Higgins, S. , Khoo, S. , Moroney, M. , O'Reilly, J. , Pilch, M. , Ryan, C. , Slattery, B. , Van Lente, E. , Walsh, E. , Walsh, J. & Hogan, V. (2018). Integrating Content Expertise and Methodological Expertise in Team-Based Settings to Address Complex Societal Issues—A Systems Perspective on Challenges. *Systems Research and Behavioral Science*. doi:10.1002/sres.2522

Hogan, M. J, Harney, O. M., & Walsh, J. C. (2017). INCluSilver: Report on Collective Intelligence Design Sessions (56 pp). *European Commission for Horizon 2020 Industrial Leadership Programme*,

Brussels. Retrieved from https://www.inclusilver.eu/wp-content/uploads/2017/09/INCluSilver_D1.5_ReportonCollectiveIntelligenceSession.pdf

Hogan, M. J. & Stein, Z. (2010). Structuring Thought: Four Unique Methods and the Challenge of Synthesis. *International Journal of Psychology Research*, 6.

Hogan, M. J. (2008). [Modest systems psychology: A neutral complement of positive psychological thinking](#). *Systems Research and Behavioural Science*, 25(6), 717-732.

Hogan, M. J., Hall, T., & Harney, O.M. (2017). Collective Intelligence Design and a New Politics of System Change. *Civitas Educationis*, 6(1), 51 – 78.

Hogan, M. J., Johnston, H., Broome, B., McMoreland, C., Walsh, J., Smale, B., Duggan, J., Andriessen, J., Leyden, K. M., Domegan, C., McHugh, P., Hogan, V., Harney, O.M. , Groarke, J., Noone, C. & Groarke, A. M. (2015). Consulting with Citizens in the Design of Wellbeing Measures and Policies: Lessons from a Systems Science Application. *Social Indicators Research*, 1-21.

Hogan, M., Harney, O., & Broome, B. (2014). Integrating Argument Mapping with Systems Thinking Tools: Advancing Applied Systems Science. In A. Okada, S.J. Buckingham Shum, & T. Sherborne (Eds.), *Knowledge Cartography* (pp. 401-421). London: Springer.

Hogan, M., Ojo, A., Harney, O. M., Ruijter, E., Meijer, A., Andriessen, J., ... & Groff, J. (2017). Governance, transparency and the collaborative design of open data collaboration platforms: understanding barriers, options, and needs. In *Government 3.0—Next Generation Government Technology Infrastructure and Services* (pp. 299-332). Cham: Springer.

Hogan, M.J. & Broome, B. (2020). Grandmothers, teachers, and teams: A commentary on the societal significance of systems thinking facilitators. *Systems Research and Behavioral Science*. <https://doi.org/10.1002/sres.2613>

Hogan, M.J., Broome, B. (2020). Facilitation and the focus on process. *Systems Research and Behavioral Science*. <https://doi.org/10.1002/sres.2639>

Jackson, M. (2019). *Critical Systems Thinking and the Management of Complexity*. NJ: John Wiley.

Thompson Long, B., Hall, T., Hogan, M., Harney, O., Doukoulos, T., & Murray, C. (2017). Using a Collective Intelligence Scenario-Based Design approach to develop a collaboration ecosystem supporting the authorship of pedagogically valuable e-books for children. *The Journal of Literacy and Technology*, 18(2).

Wallace, E., Hogan, M., Noone, C., & Groarke, J. (2018). Investigating components and causes of sabotage by academics using collective intelligence analysis. *Studies in Higher Education*, 1-19.

Warfield, J., & Cardenas, R. (1994). *A handbook of interactive management*. Ames: Iowa State University Press.